CSC6730 – ADVANCED NETWORKING

A 30,000 FOOT OVERVIEW OF THE INTERNET

Susmit Shannigrahi
Computer Science
sshannigrahi@tntech.edu

Tennessee TECH

Topics covered in this Lecture

- Where the Internet came from
- Design decisions
- Mechanisms that drive it
- Problems
- Where do we go from here?

Early days of the Internet

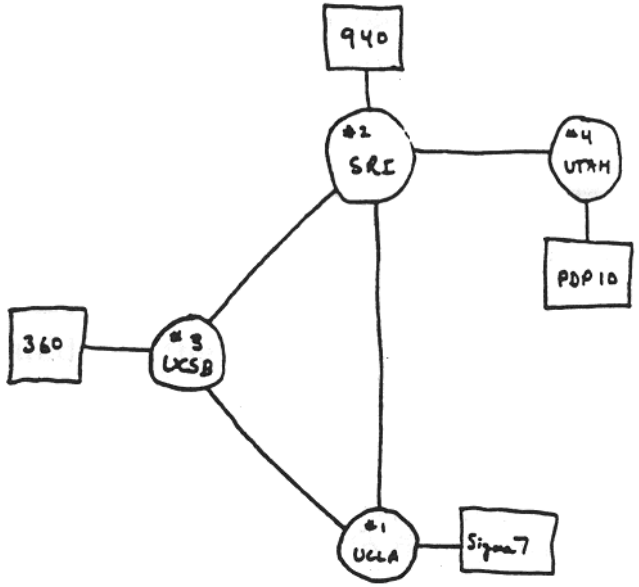

FIGURE 6.2 Drawing of 4 Node Network
(Courtesy of Alex McKenzie)

Fast forward a few years - 2021

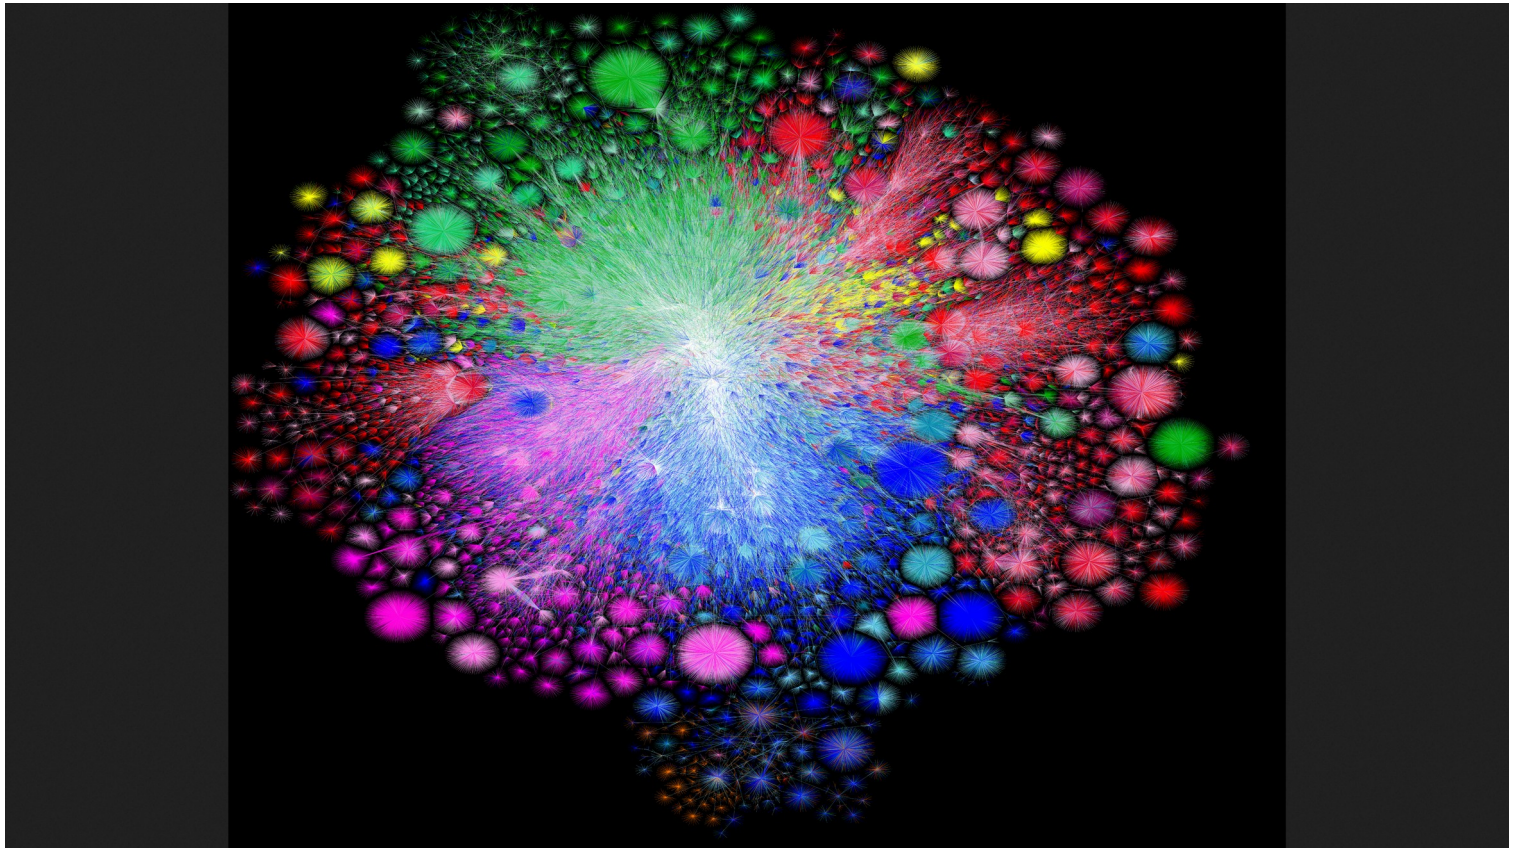

Each blob is an autonomous system

Colors represent geographical region

http://renderbot.nyc1.opte.org/20210116.1200.attempt_1/attempt_1_run_20210116.1200.coords_10000x8000_dark.png

What does the infrastructure look like?

https://
www.submarinecablemap.com/

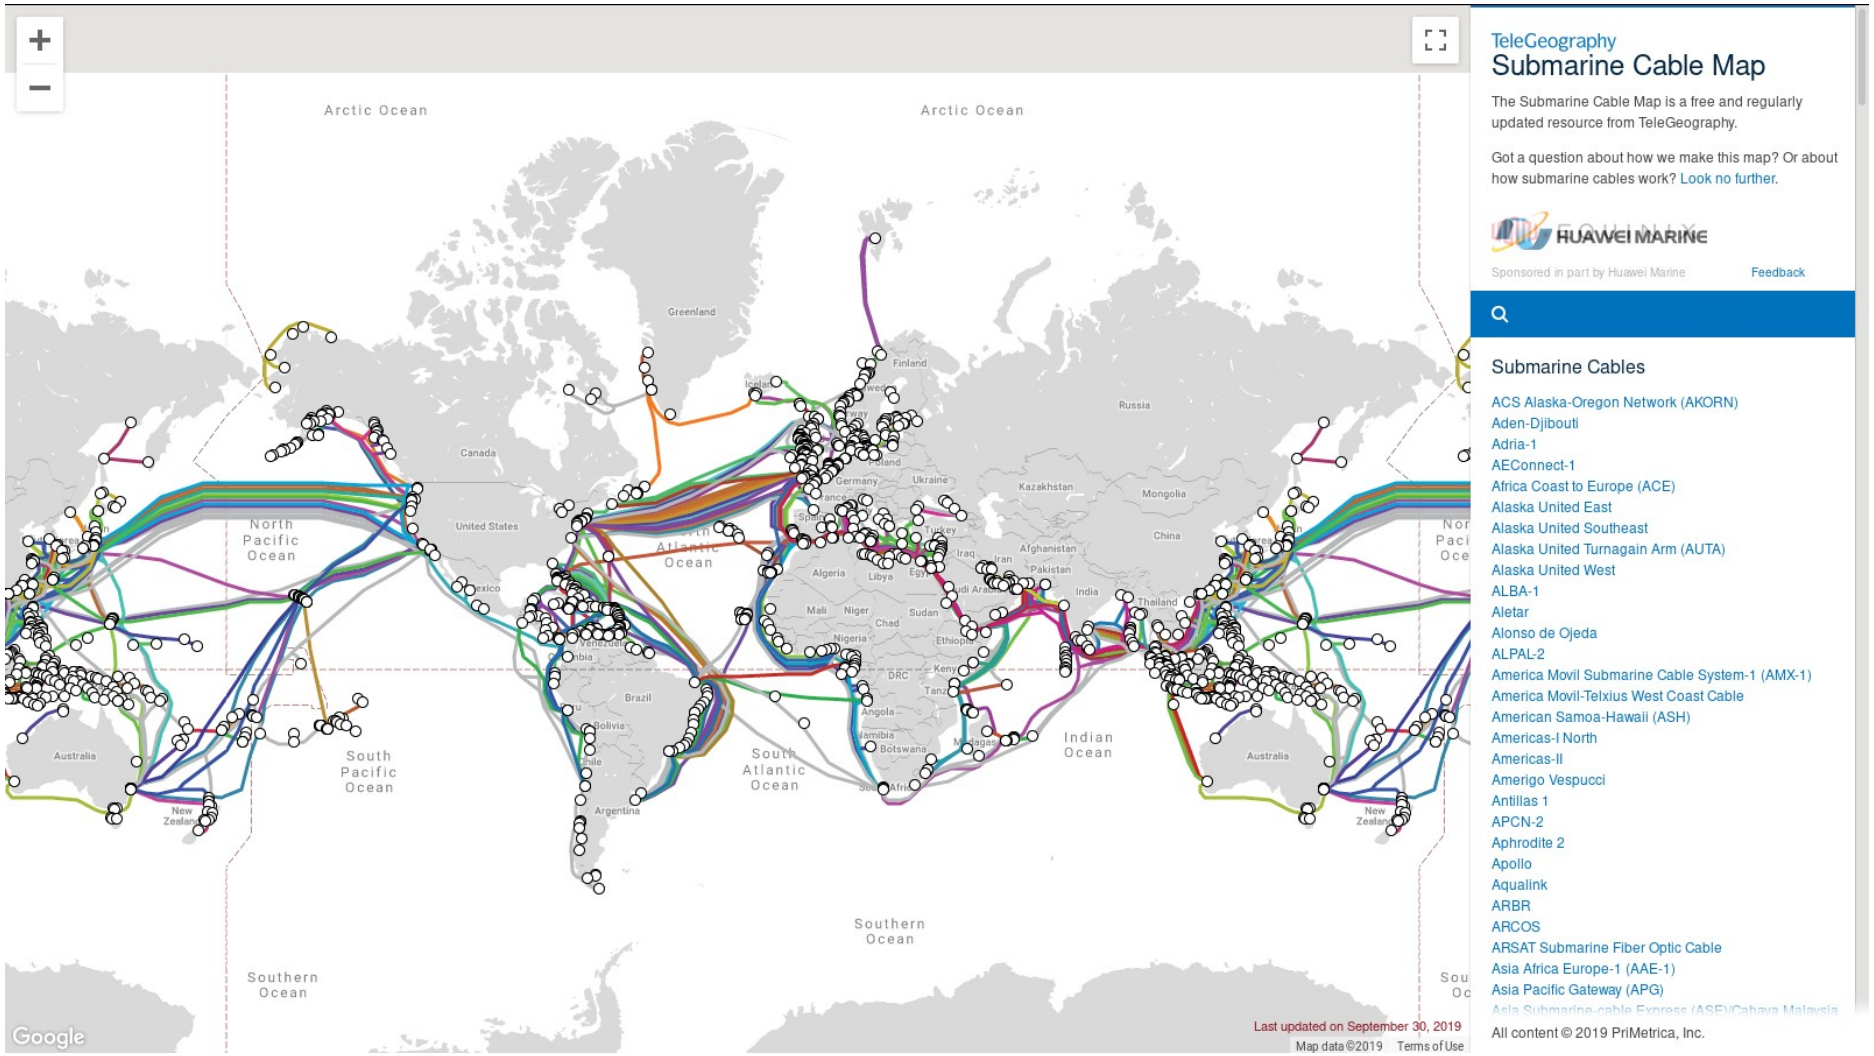

Number of hosts on the Internet in 2018

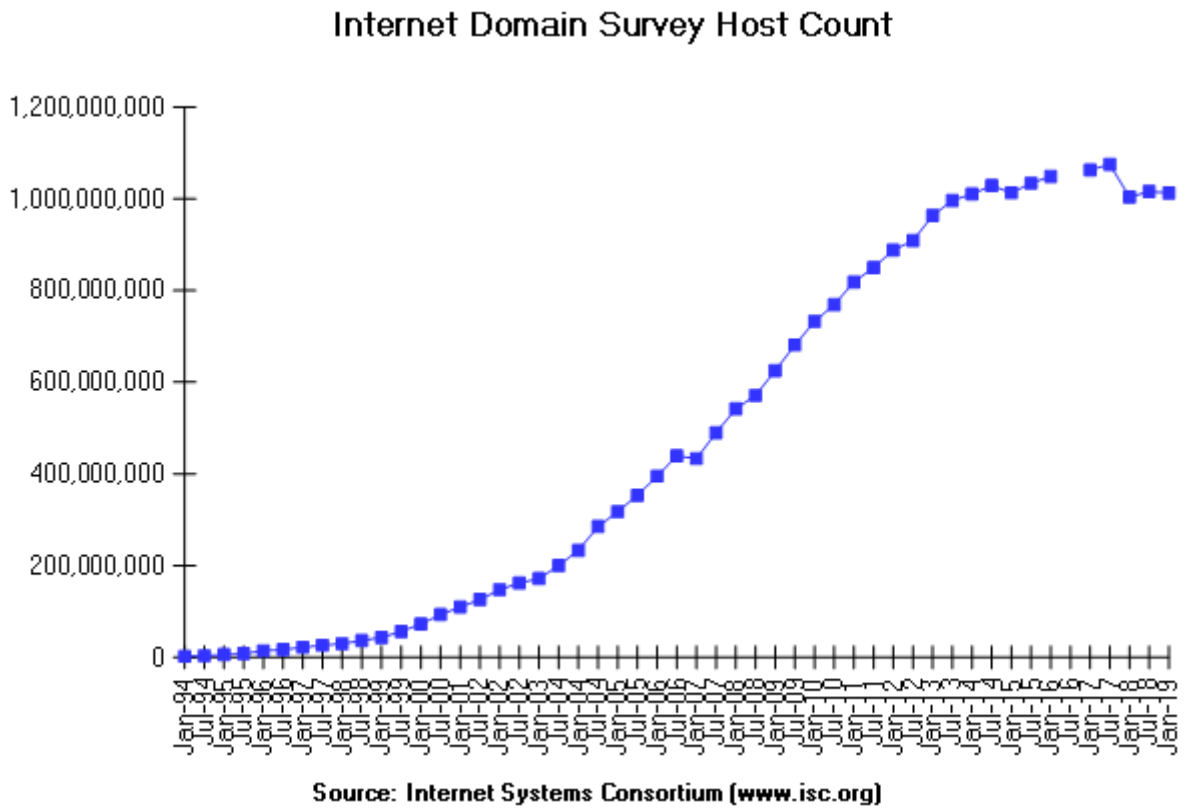

What is missing from this?

How many Autonomous Systems?

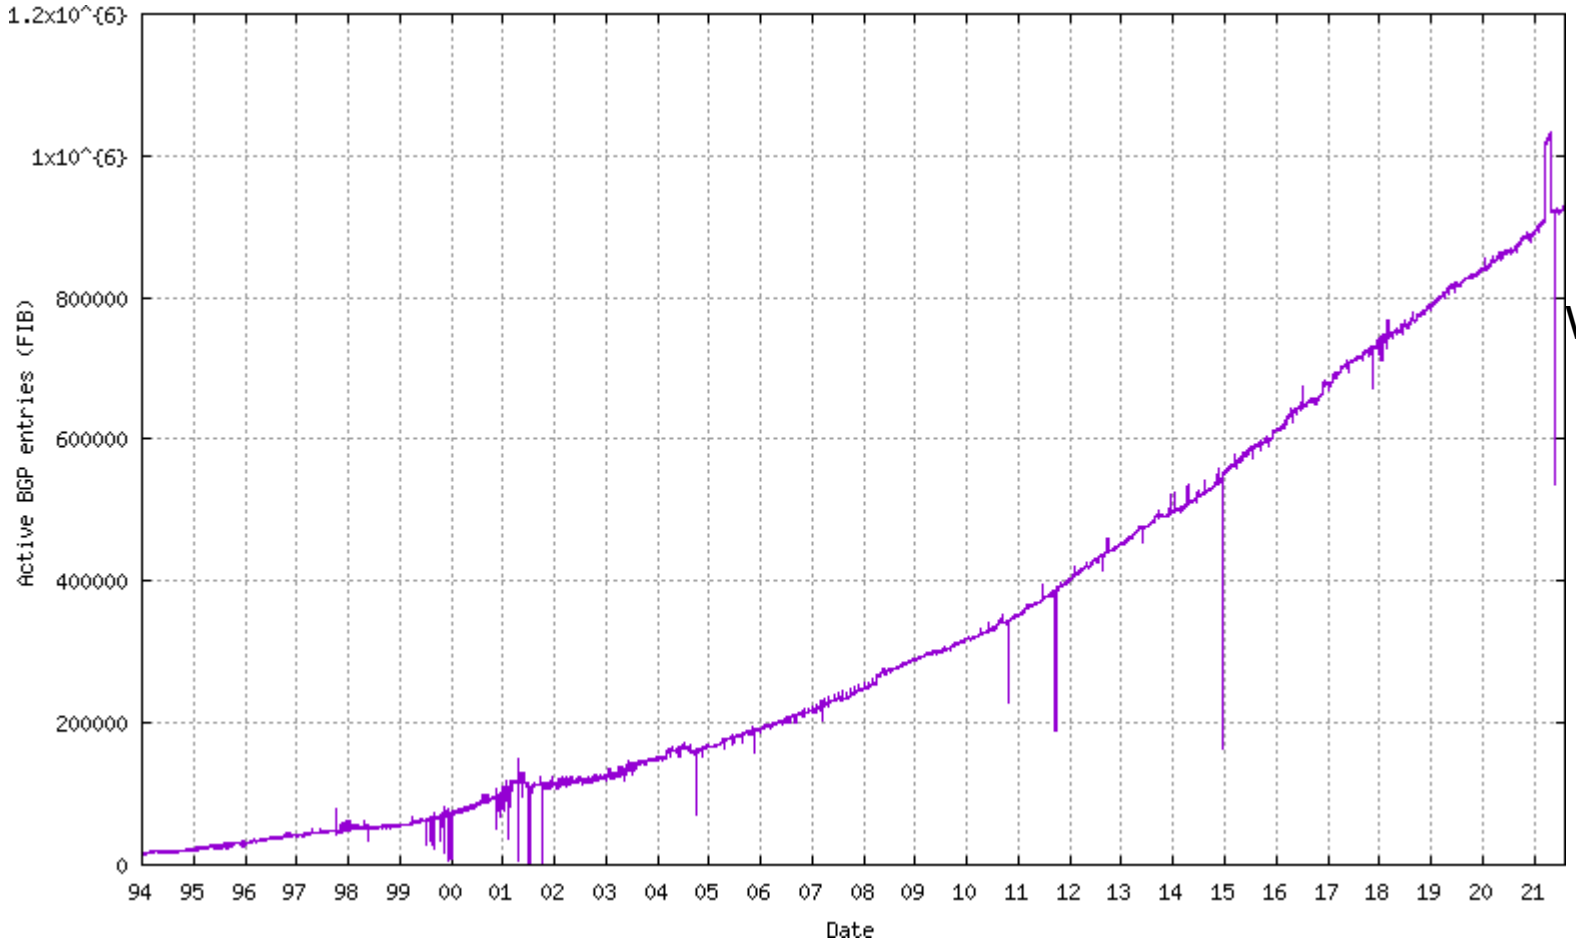

What is missing from this?

How many Users?

Share of the population using the Internet, 2017

All individuals who have used the Internet in the last 3 months are counted as Internet users. The Internet can be used via a computer, mobile phone, personal digital assistant, games machine, digital TV etc.

Our World in Data

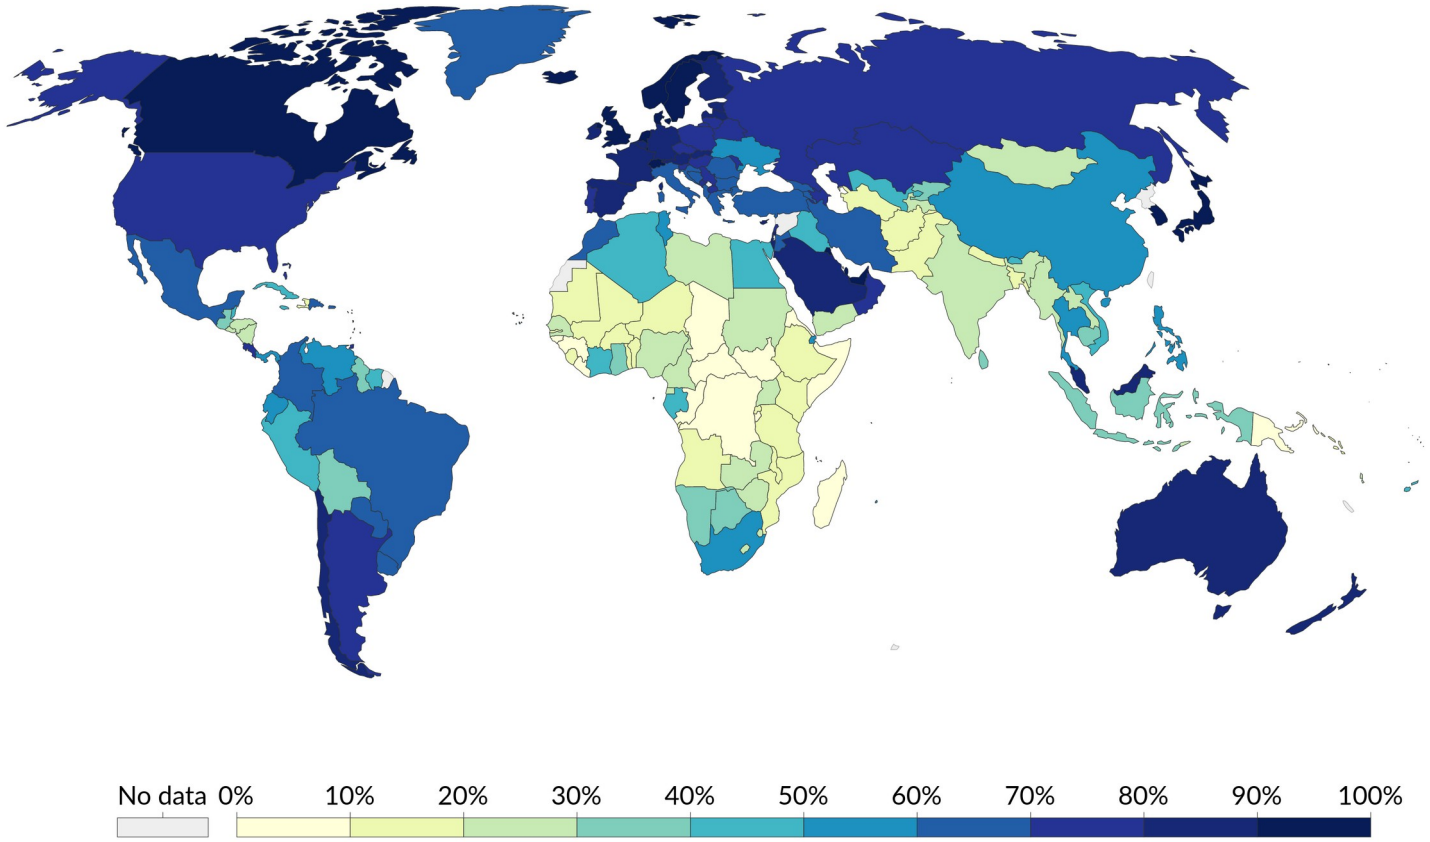

Source: World Bank

OurWorldInData.org/technology-adoption/ • CC BY

Topics covered in this Lecture

- Where the Internet came from
- Design decisions
- Mechanisms that drive it
- Problems
- Where do we go from here?

Bob Kahn’s 4 ground rules

- Each distinct network would have to stand on its own and no internal changes could be required to any such network to connect it to the Internet.
  - Network of Networks!

- Communications would be on a best effort basis. If a packet didn't make it to the final destination, it would shortly be retransmitted from the source.
  - Collisions!

- Black boxes would be used to connect the networks; these would later be called gateways and routers. There would be no information retained by the gateways about the individual flows of packets passing through them, thereby keeping them simple and avoiding complicated adaptation and recovery from various failure modes.
  - Intelligence at the end-points not in the network

- There would be no global control at the operations level.

https://arxiv.org/html/cs/9901011

Layering

- Common and stable interfaces between layers
- Internal changes do not affect the other layers

Open Design

- Open standards
- Open protocols
- Open Code!
- Has greatly helped with the Internet’s spread!

Topics covered in this Lecture

- Where the Internet came from
- Design decisions
- Mechanisms that drive it
- Problems
- Where do we go from here?

What do the Internet layers look like?

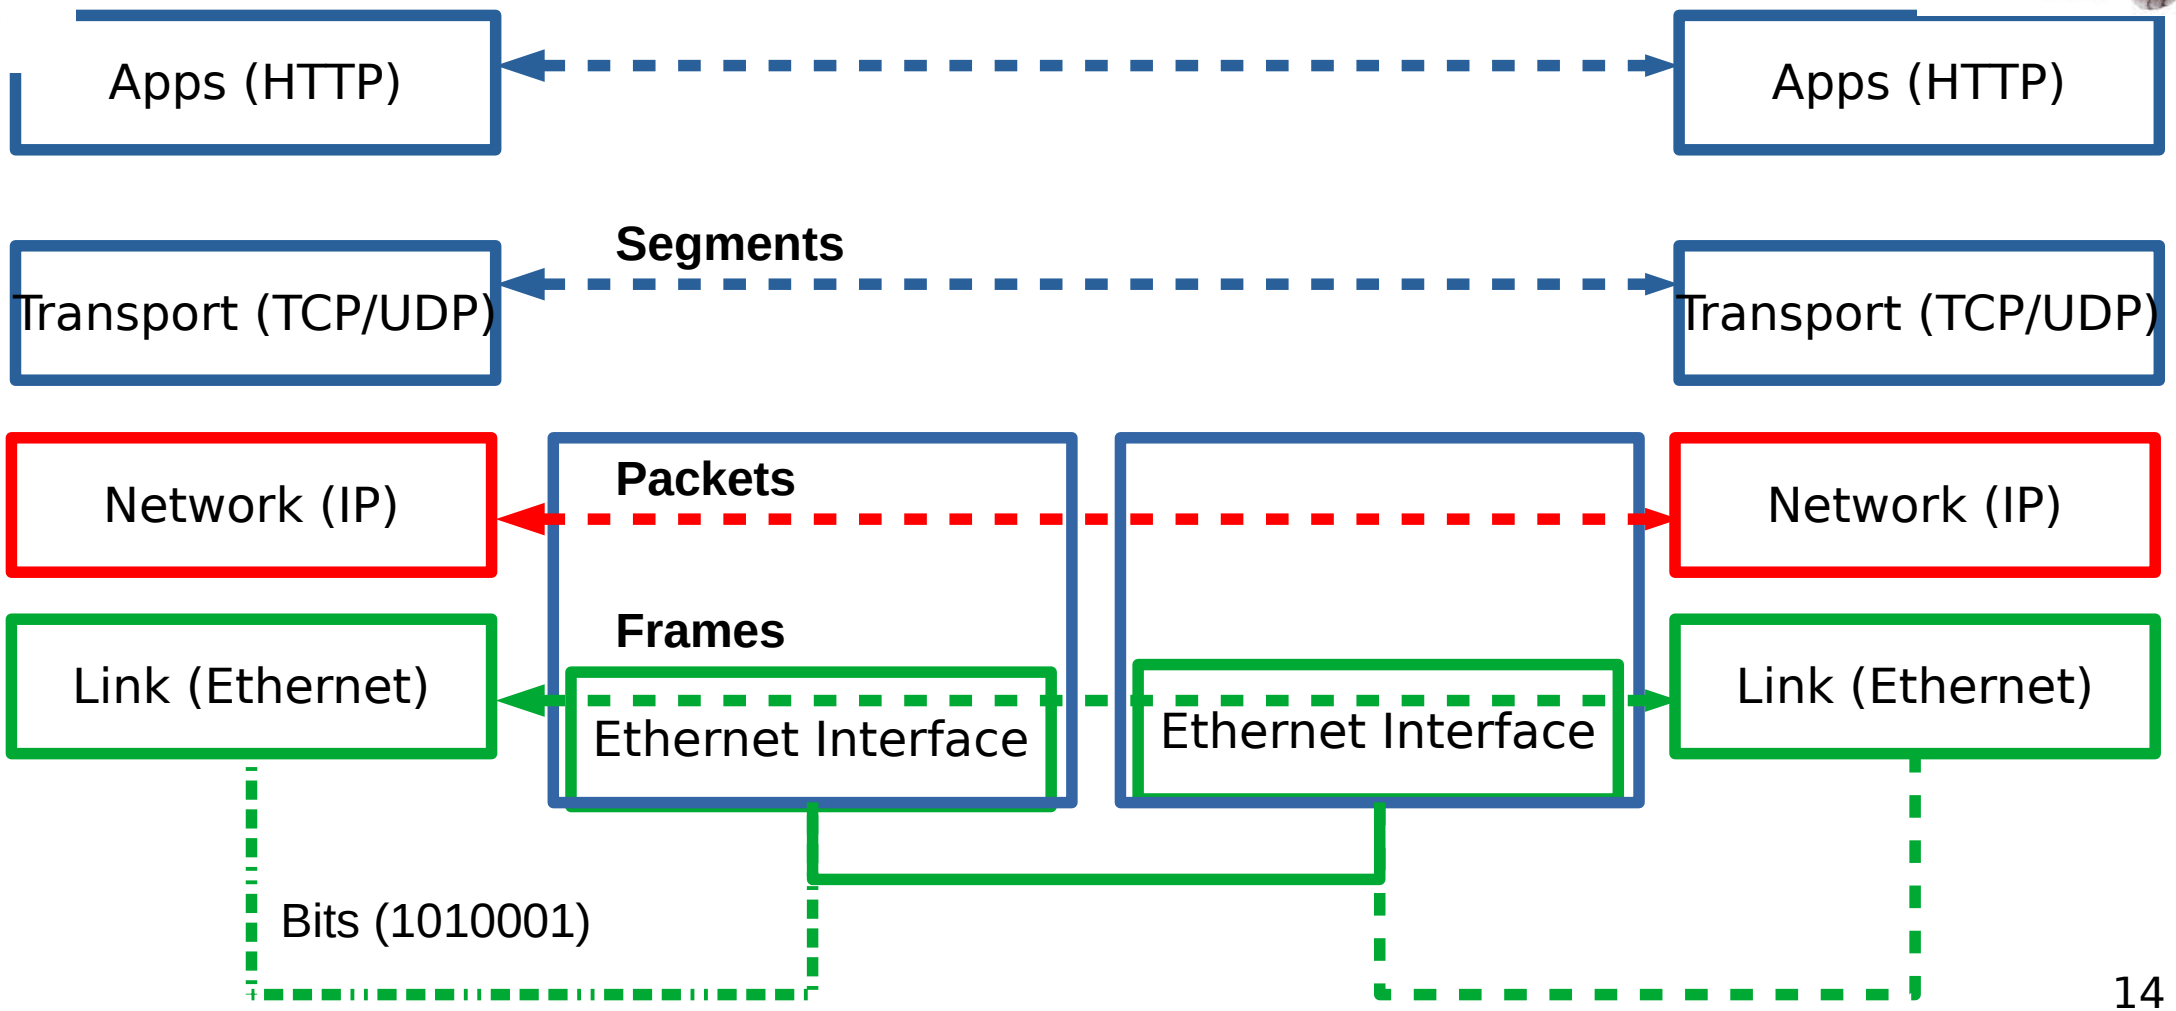

Link Layer!

https://
www.submarinecablemap.com/

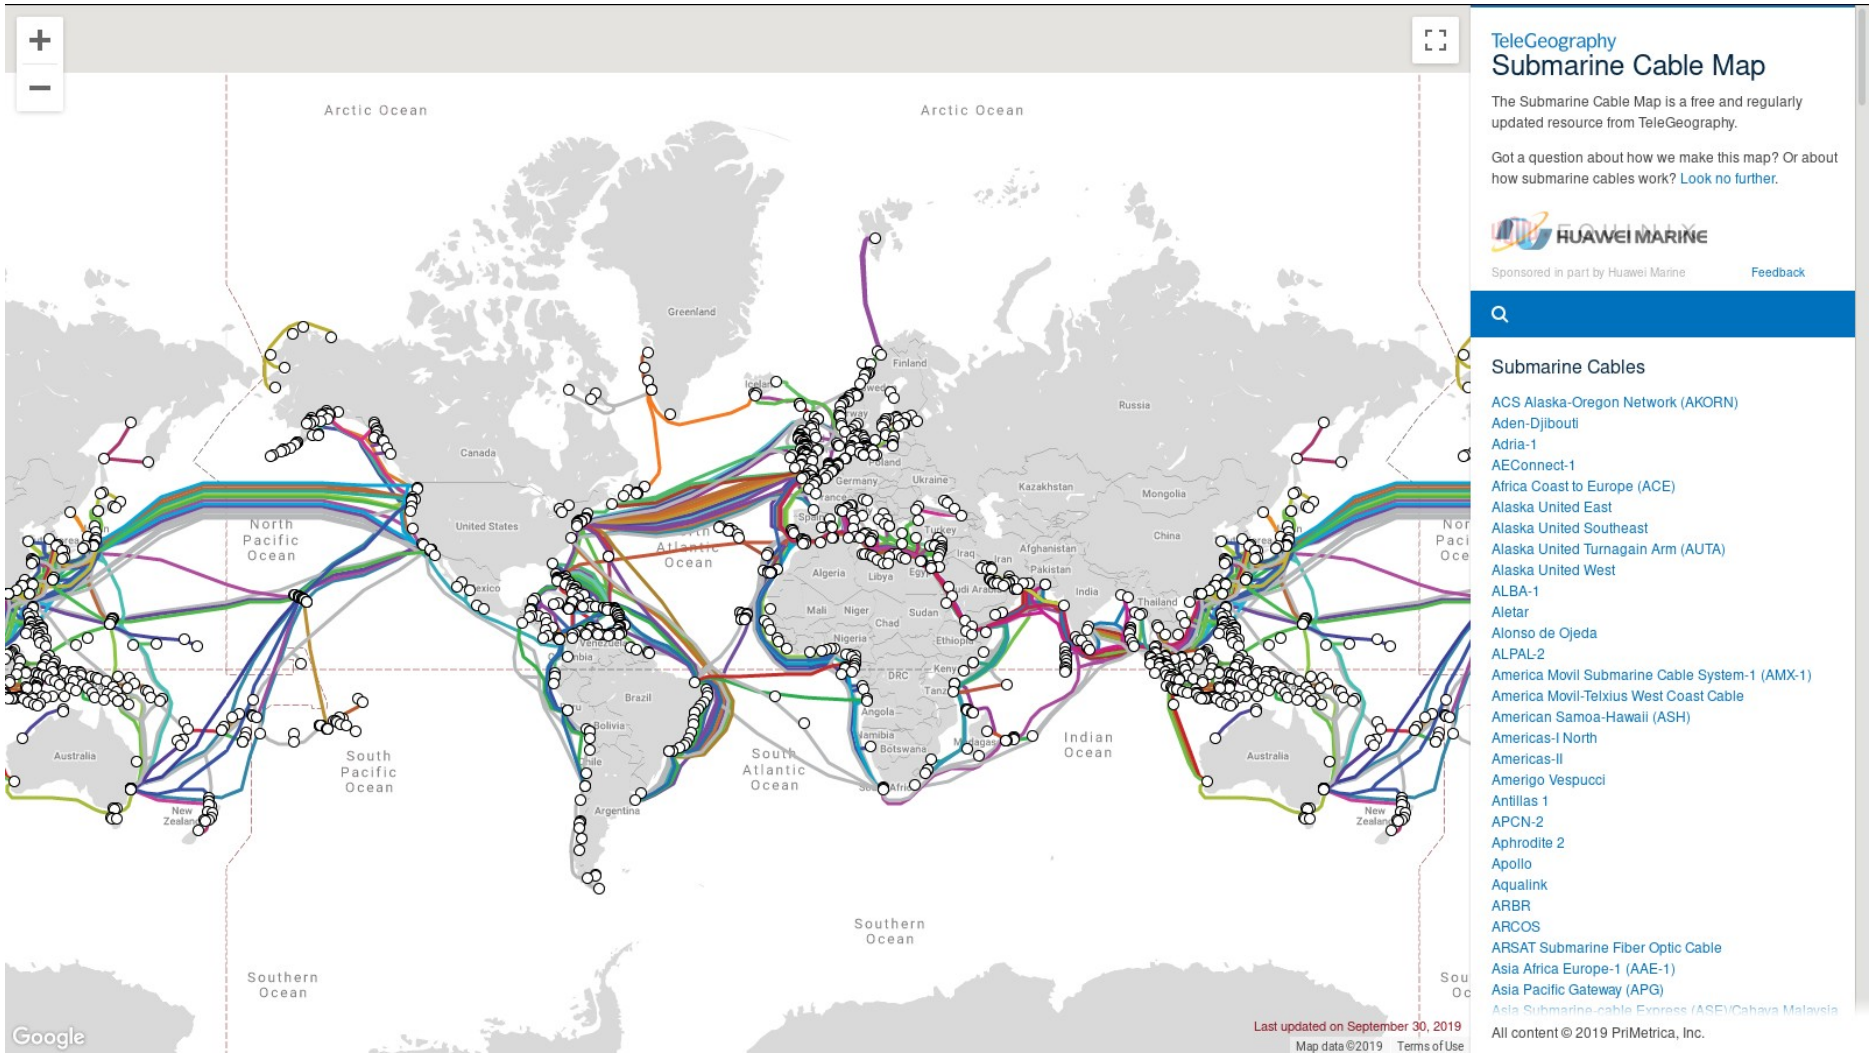

Switching

- Switch
  - A mechanism to interconnect links to form a large network
  - Forward **frames**
  - Separate the collision domains
  - Filter packets between LANs
  - Connects two or more LAN segments - **Bridging**

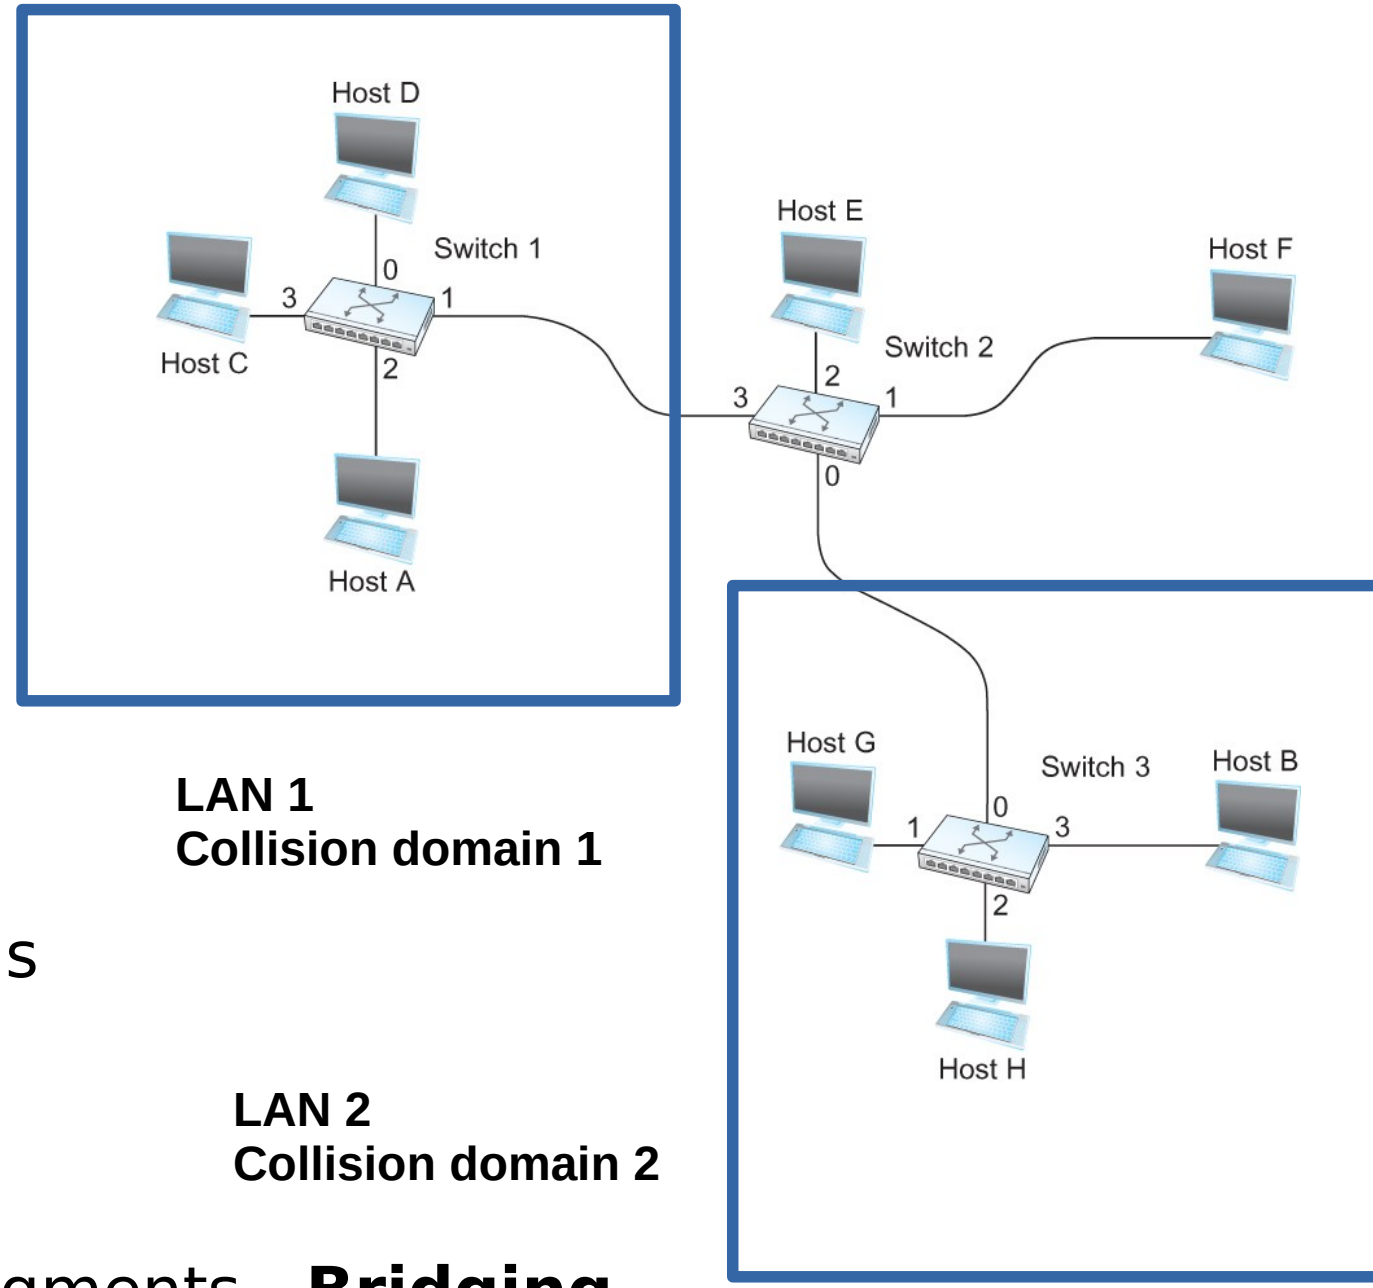

Switching Table

- Unknown destination → send out on all Interfaces **(flooding)**
  - **Skip the incoming interface**

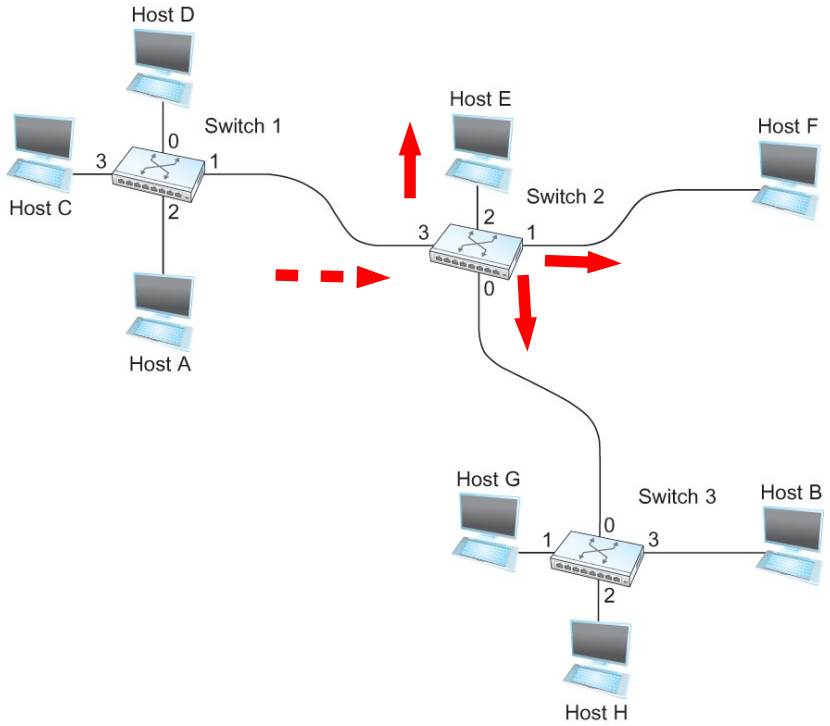

Destination, Port	
	--
A	3
B	0
C	3
D	3
E	2
F	1
G	0
H	0

Forwarding Table for Switch 2

Switches are self learning!

- Inspect the source MAC address
  - **What is a mac address?**
- Associate mac address and incoming interface
- Store this association for later use, (for some time)
  - aging-timer

Switches are self learning!

- Inspect the source MAC address
  - **What is a mac address?**
- Associate mac address and incoming interface
- Store this association for later use, (for some time)
  - aging-timer

Internet Protocol (IP)

- What is an internetwork?
  - An arbitrary collection of networks interconnected to provide some sort of host-host to packet delivery service
  - Note the different types of network

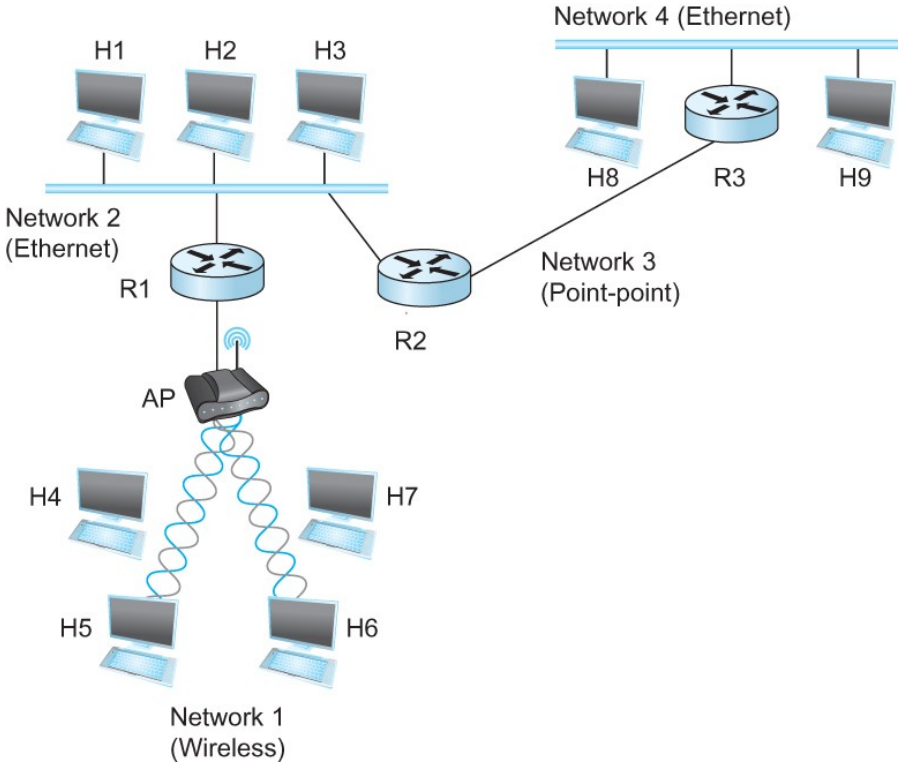

Internet Protocol (IP)

- What is an internetwork?
  - An arbitrary collection of networks interconnected to provide some sort of host-host to packet delivery service
  - Note the different types of network

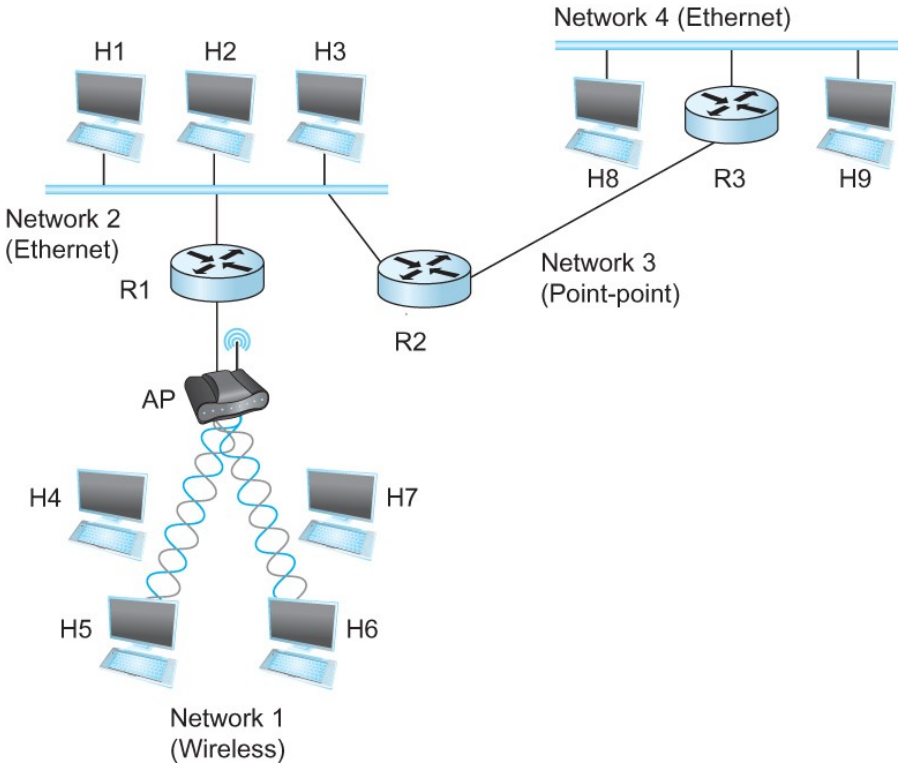

Back to Addressing

- A 32 bit number in quad-dot notation
- Identifies an ***Interface***
  - ***A host might have several interfaces!!!***
- *129.82.138.254*

  10000001.01010010.10001010.111

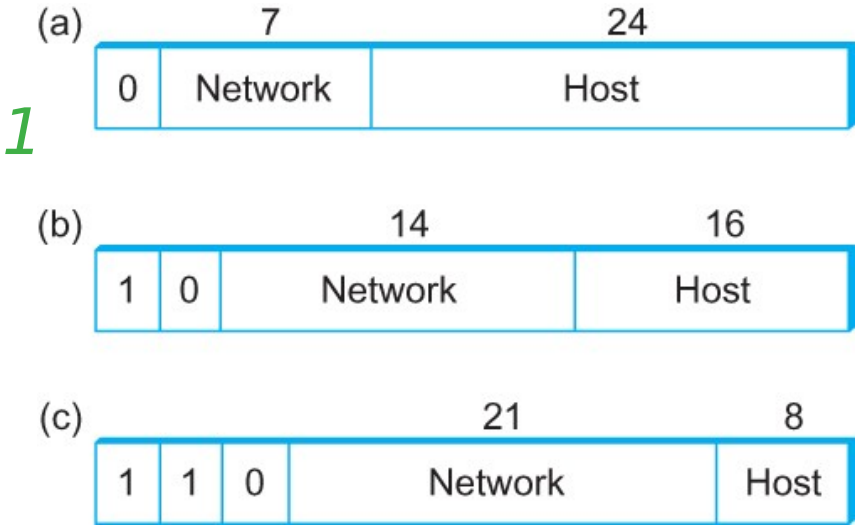

Subnets Revisited

Recipe:

- Create isolated networks – *subnets*

- No longer need to know individual Ips – knowing the subnet is enough
  - *223.1.1.0/14 → Interface 2*

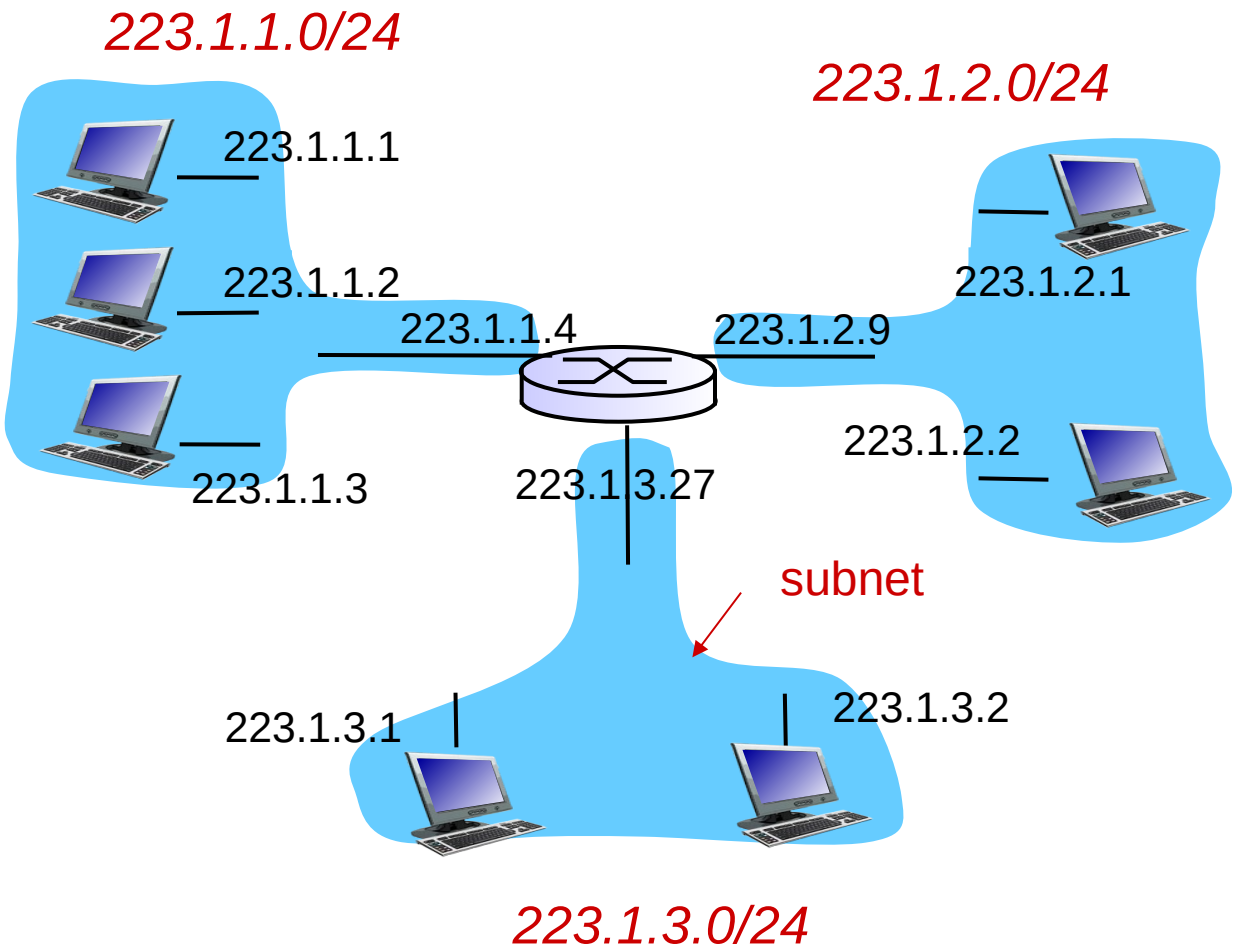

subnet mask: /24

Subnets (Prefixes) scales the Internet

- Addresses are allocated in contiguous prefixes (tntech 149.149.0.0/16)
- Routing protocols operate based on prefixes (how do I reach 149.149.0.0/16)?

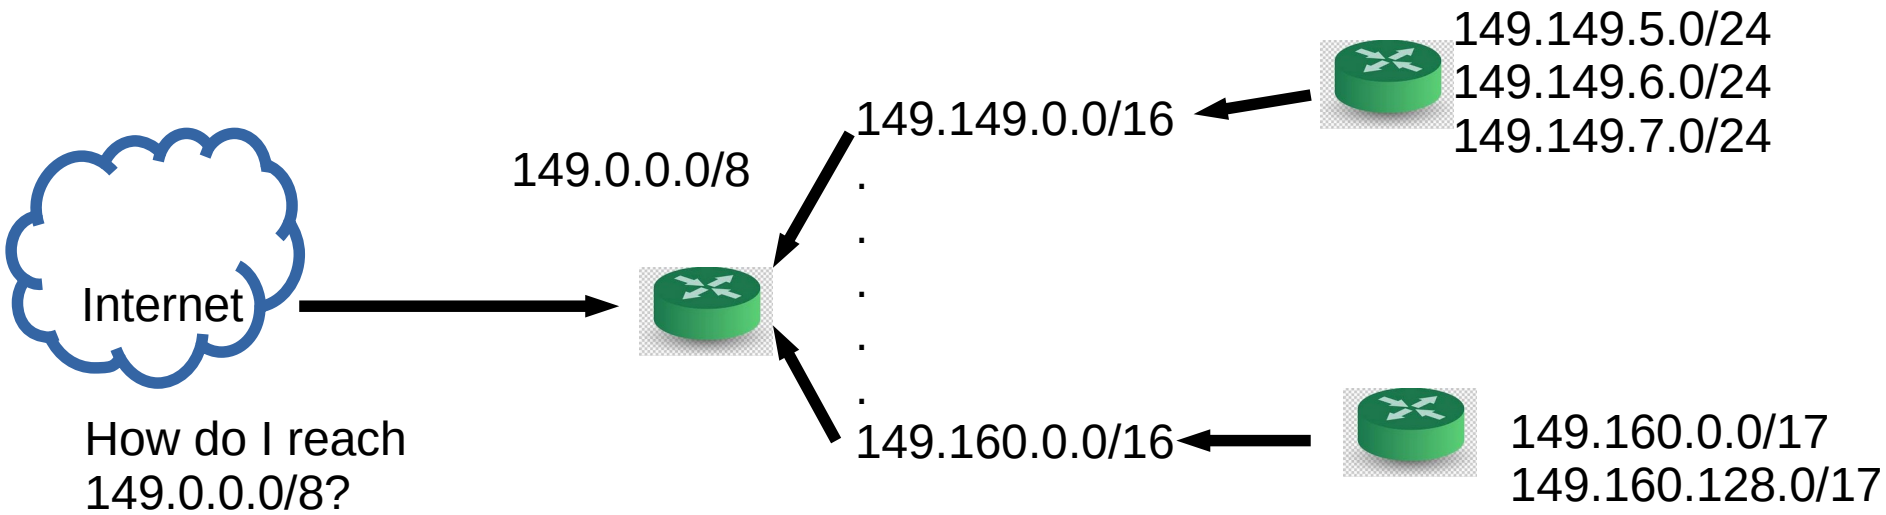

Not

How do I reach 149.149.5.0/24

How do I reach 149.149.6.0/24

Who gets what prefix?

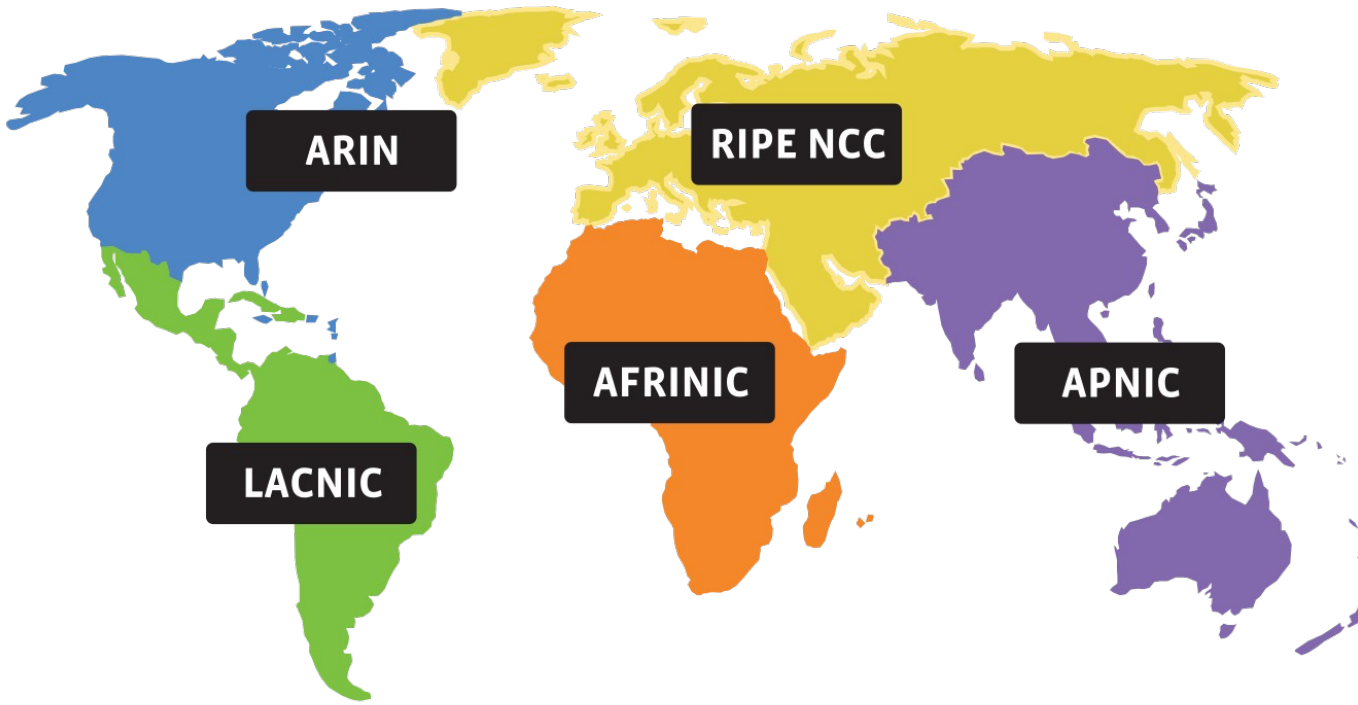

0. Internet Corporation for Assigned Names and Numbers (ICANN) – Decides which RIRs get what address

1. Regional Internet Registries (RIRs) – Which orgs get what address

2. ISPs – Which customers get which address

IPv4 Crisis

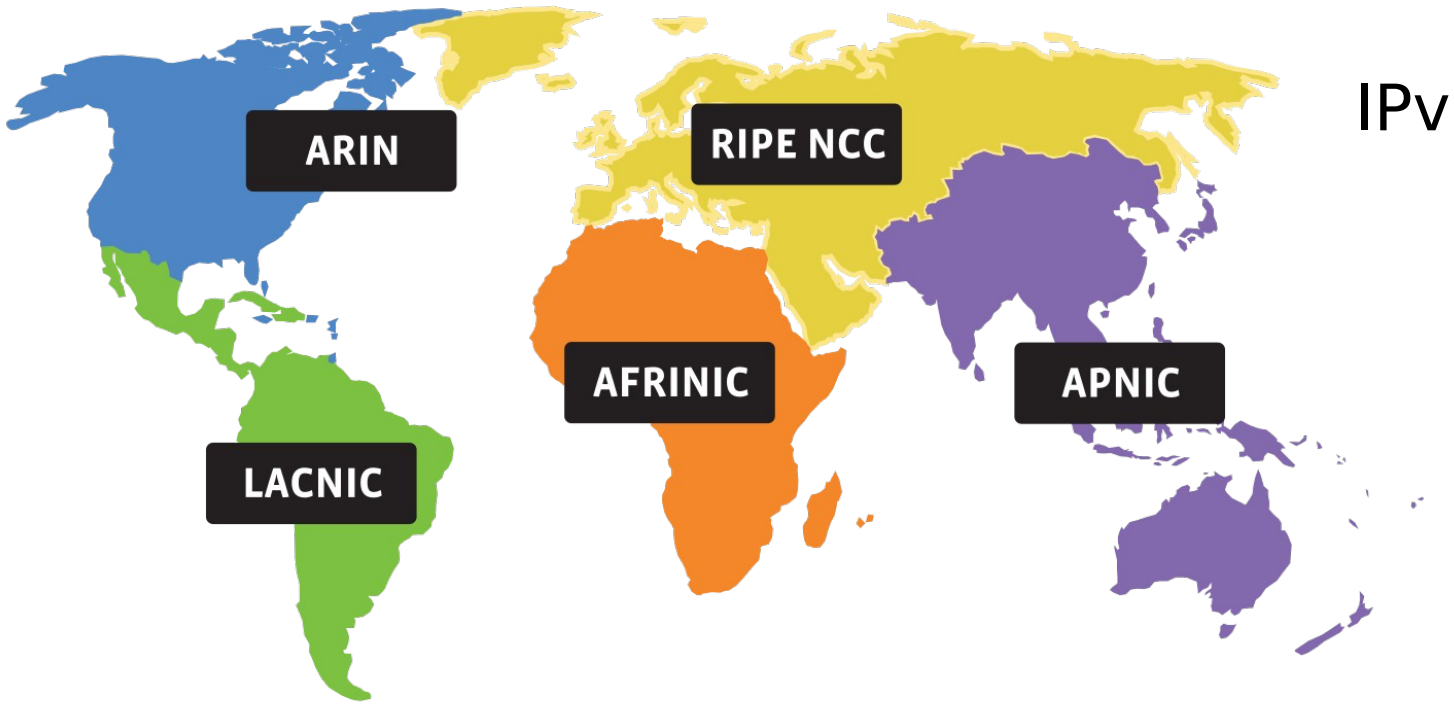

We have officially run out of IPv4 addresses

IPv6 to the rescue – use it.

How do we talk to the other networks?

Some sort of routing table

- Like your GPS

- Allows you to create paths to other networks

Forwarding vs Routing

- Forwarding:
  - to select an output port based on destination address and routing table
  - **Local path**

SubnetNumber	SubnetMask	NextHop
128.96.34.0	255.255.255.128	Interface 0
128.96.34.128	255.255.255.128	Interface 1
128.96.33.0	255.255.255.0	R2

- Routing:
  - process by which routing table is built
  - **End-to-end path**

Why bother?

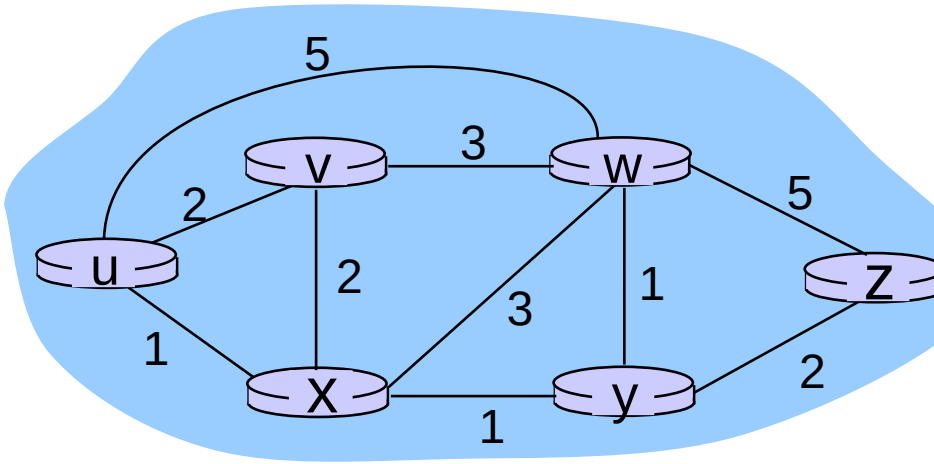

- Quality of path affects performance
  - Longer path = more delay
- Balance path usage, avoid congested paths
- Deal with failures

SubnetNumber	SubnetMask	NextHop
128.96.34.0	255.255.255.128	Interface 0
128.96.34.128	255.255.255.128	Interface 1
128.96.33.0	255.255.255.0	R2

Router architecture overview

Two key router functions:
- run routing algorithms/protocol (RIP, OSPF, BGP)
- *forwarding* datagrams from incoming to outgoing link

Control Plane = routing
Vs
Data Plane = forwarding

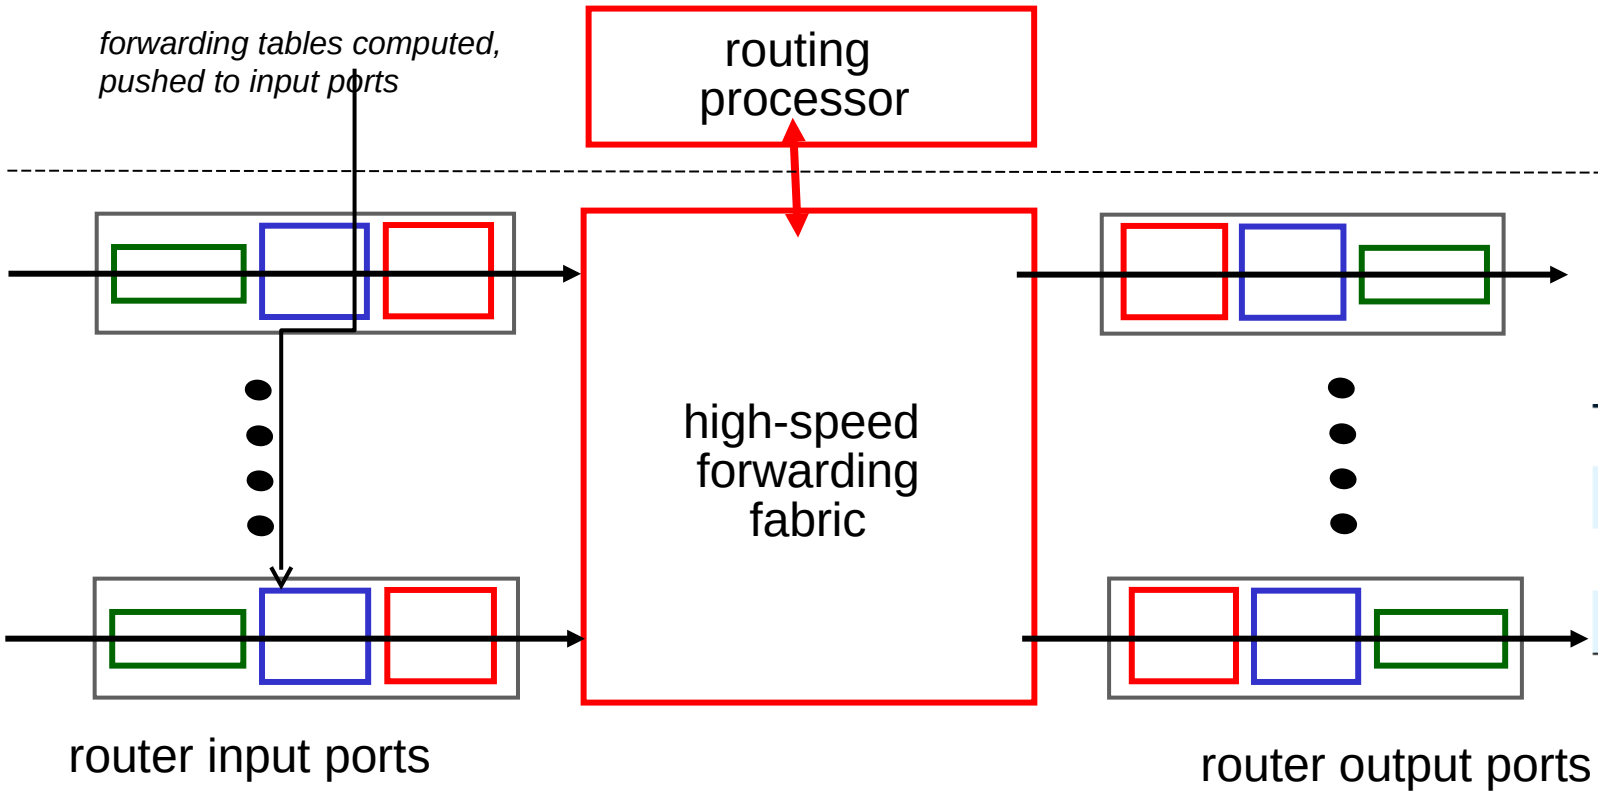

SubnetNumber	SubnetMask	NextHop
128.96.34.0	255.255.255.128	Interface 0
128.96.34.128	255.255.255.128	Interface 1
128.96.33.0	255.255.255.0	R2

Router architecture overview

Two key router functions:

- run routing algorithms/protocol (RIP, OSPF, BGP)
- *forwarding* datagrams from incoming to outgoing link

Control Plane = routing
Vs
Data Plane = forwarding

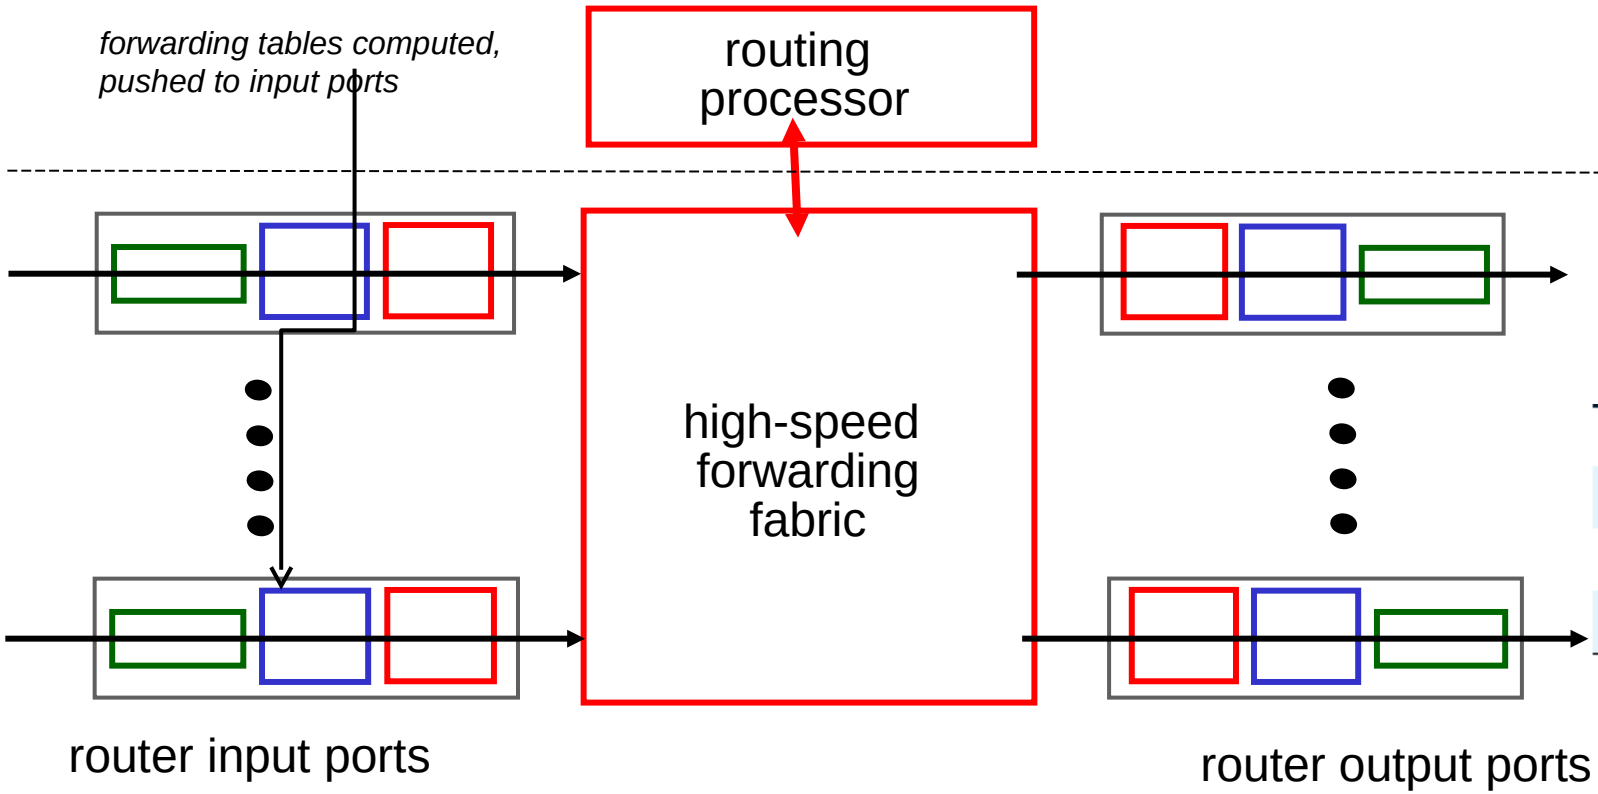

SubnetNumber	SubnetMask	NextHop
128.96.34.0	255.255.255.128	Interface 0
128.96.34.128	255.255.255.128	Interface 1
128.96.33.0	255.255.255.0	R2

Scaling Routing!
Internet in the 1990s

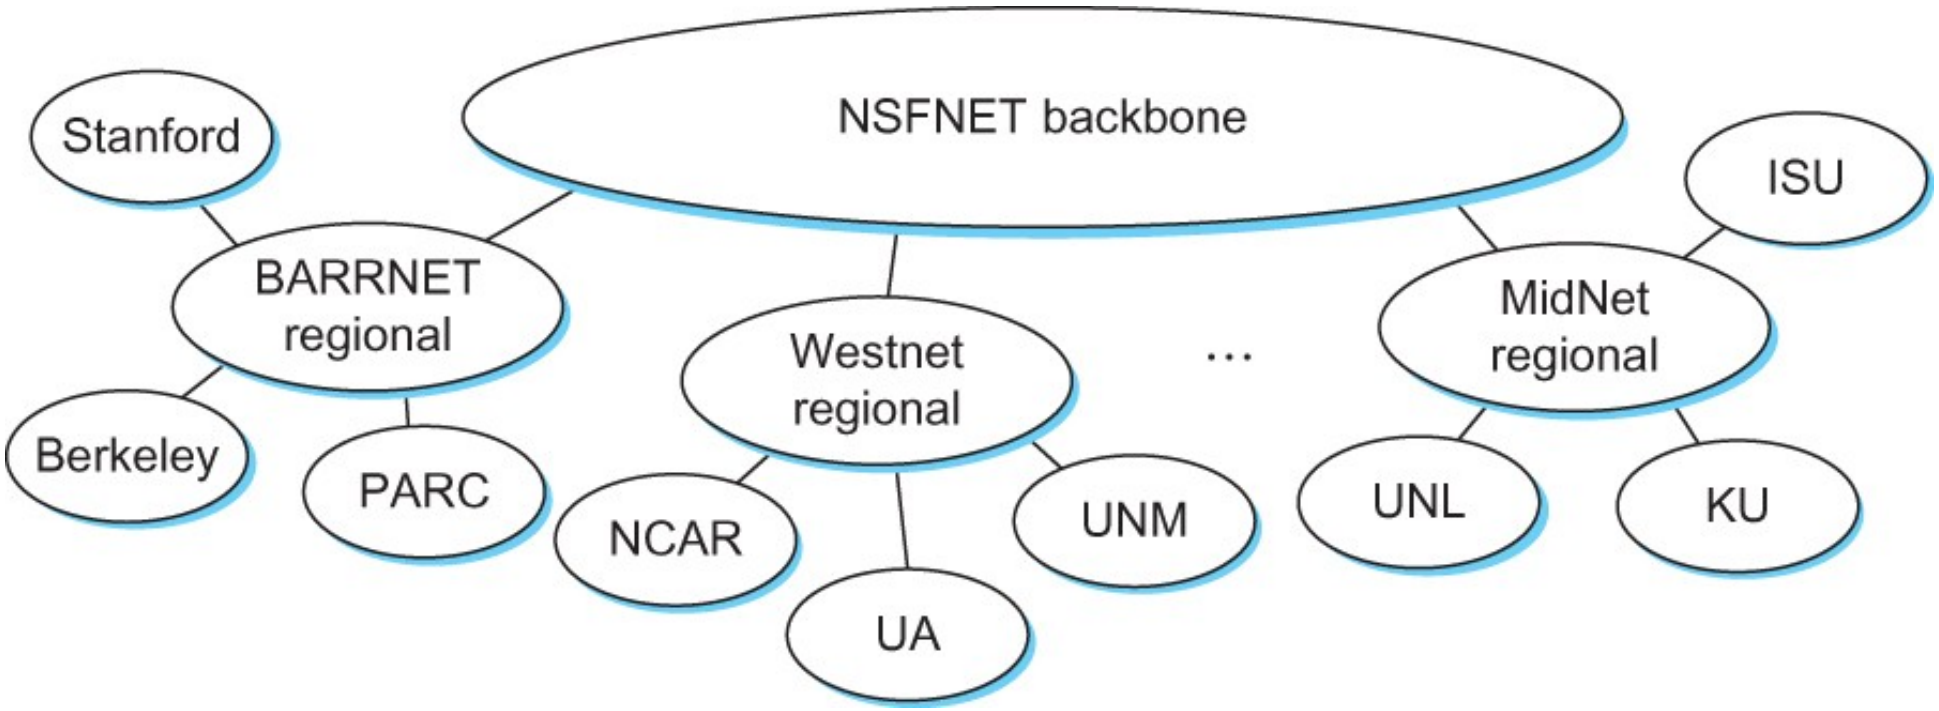

Internet now

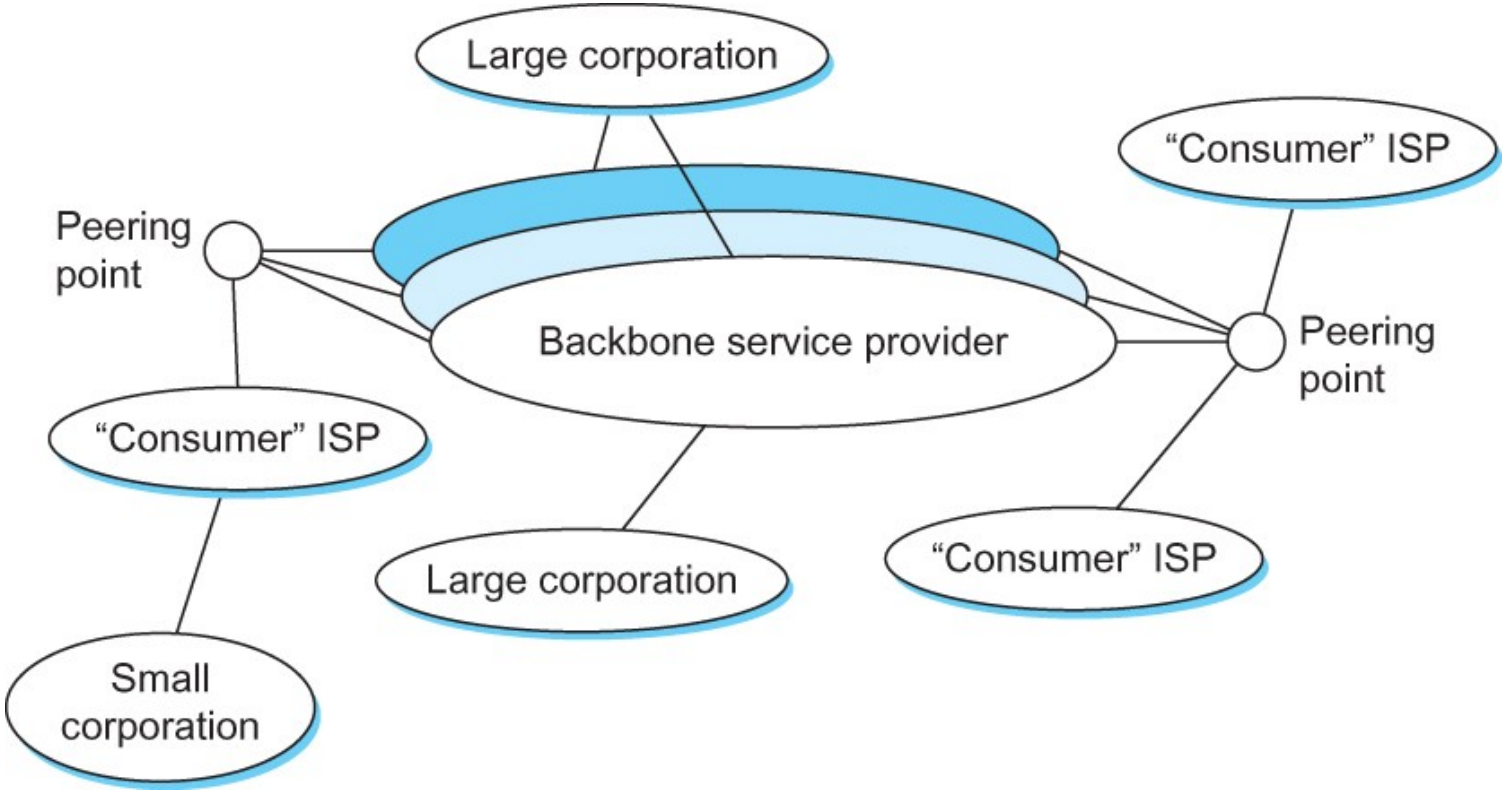

Hierarchical routing - Policy

scale: with 600 million destinations:

- can’t store all dest’s in routing tables!
- routing table exchange would swamp links!

administrative autonomy

- internet = network of networks
- each network admin may want to control routing in its own network

Autonomous systems (ASes)

- AS
  - A set of routers under a single technical administration
  - What happens inside an AS stays within that AS!
    - That is, AS decides routing metrics internally

BGP-4: Border Gateway Protocol

- Assumes the Internet is an arbitrarily interconnected set of AS's.

- Local traffic – within the AS
- Transit traffic – from AS1 to AS3 via AS2
- Three types of AS’s
  - *Stub AS*
  - *Multihomed AS*
  - *Transit AS*

BGP-4: Border Gateway Protocol

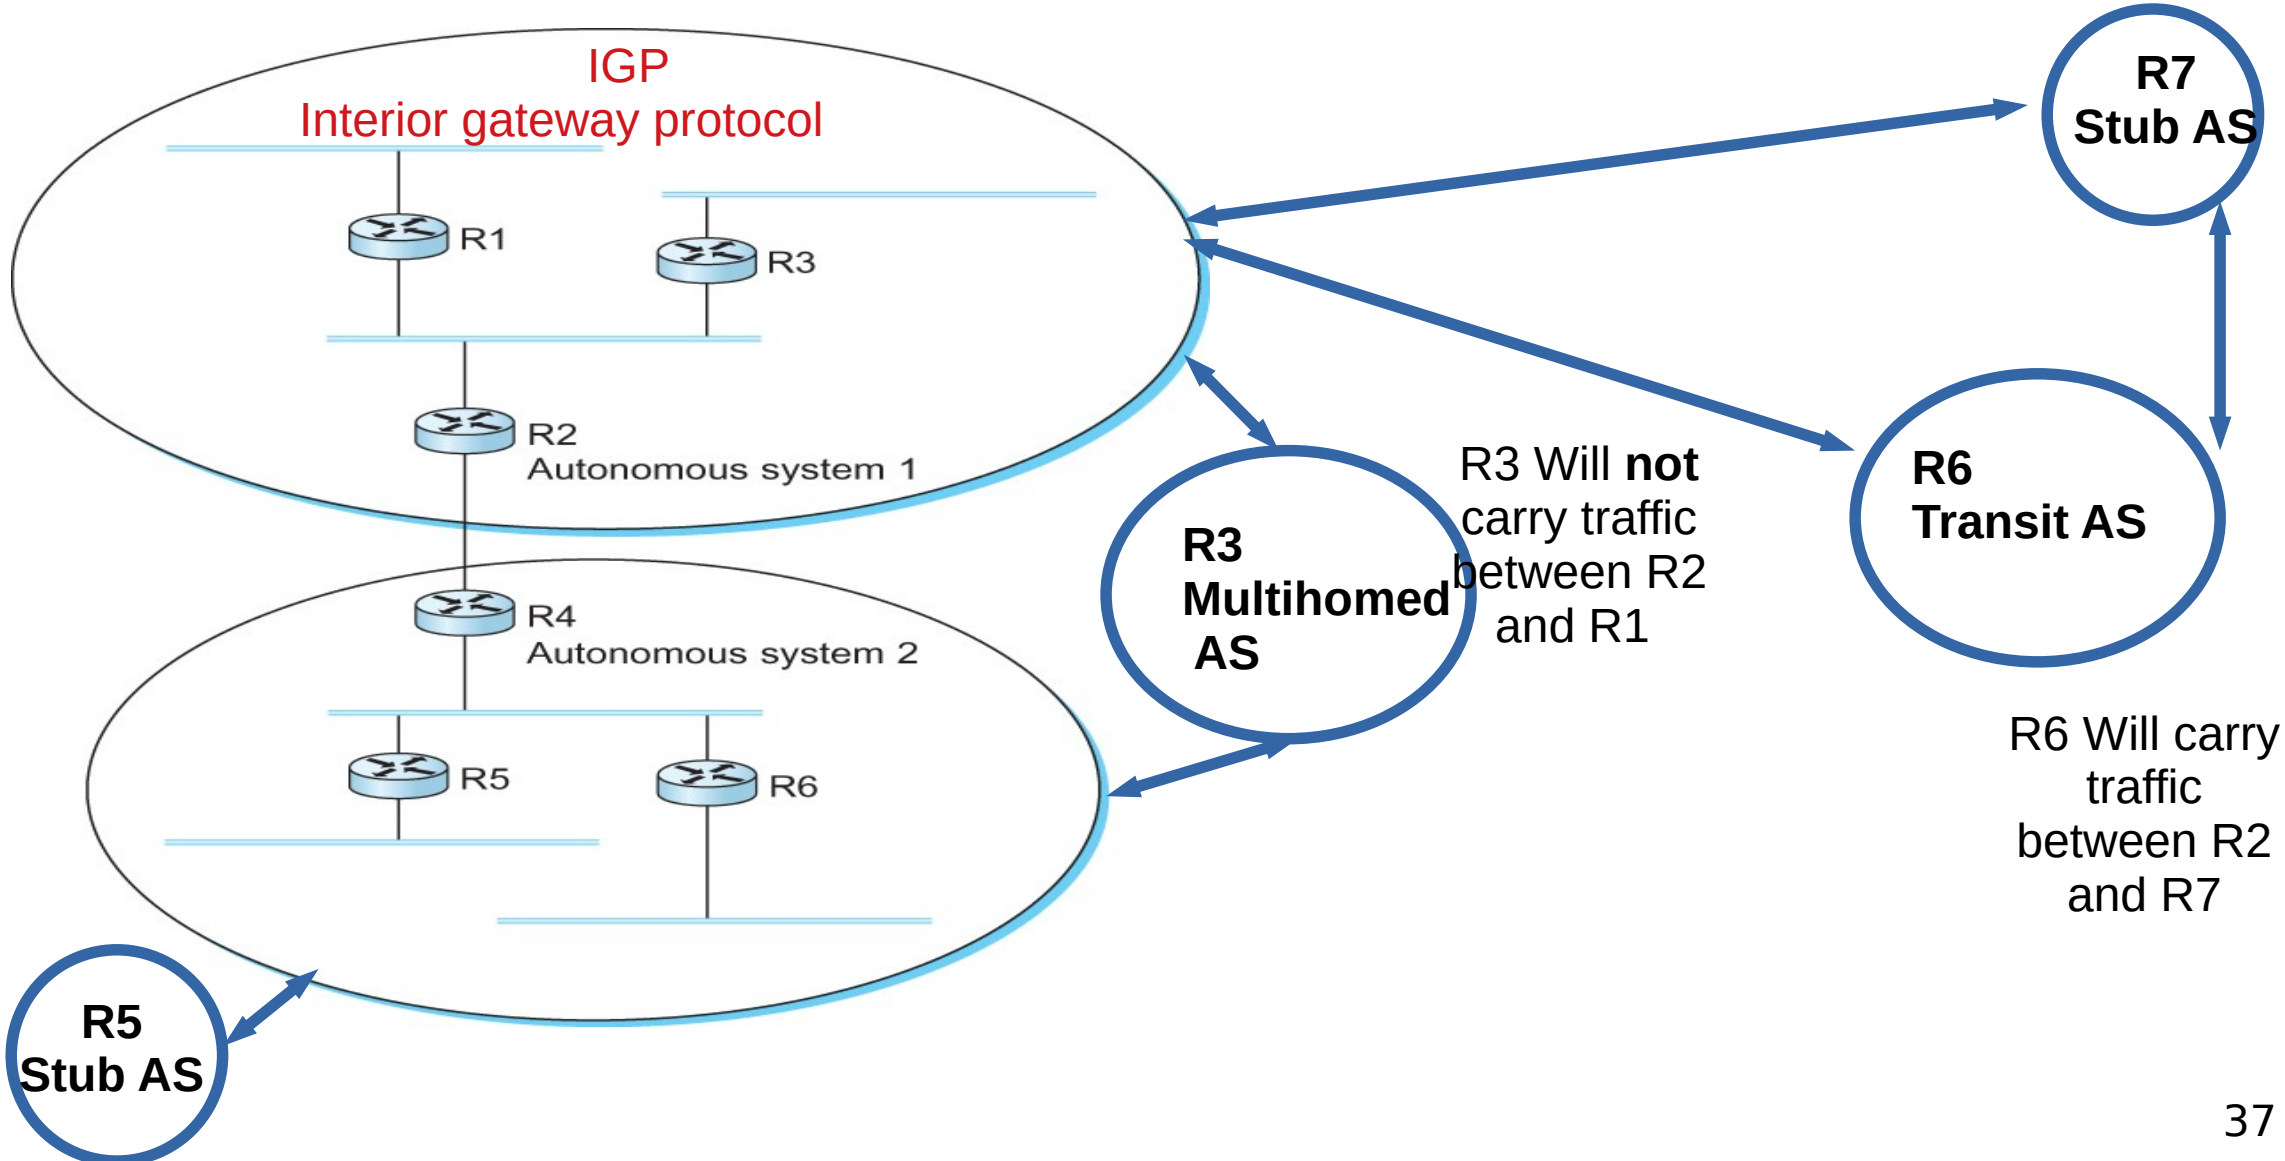

BGP - goals

- The goal of Inter-domain routing is to find **any path** to the intended destination that is **loop free**
  - **We are concerned with reachability than optimality**
  - Finding path anywhere close to optimal is considered to be a great achievement
- Why?

BGP - Goals

- Scalability: Forward any packet destined anywhere in the Internet
  - Having a routing table that will provide a match for any valid IP address
- Autonomous nature of the domains
  - impossible to calculate meaningful costs for a path crossing multiple ASs
  - A cost of 1000 is great at provider 1, terrible at provider 2
- Issues of trust
  - Provider A might be unwilling to believe certain advertisements from provider B

BGP: Path vector protocol

- Send the whole path with the routing update
- Loops are detected if an AS finds itself in the path
  - Reject if so
  - Accept otherwise
- Add self to the path and advertise to the neighbors
- Advantage: No loops, Local decision before advertising

BGP: Path vector protocol

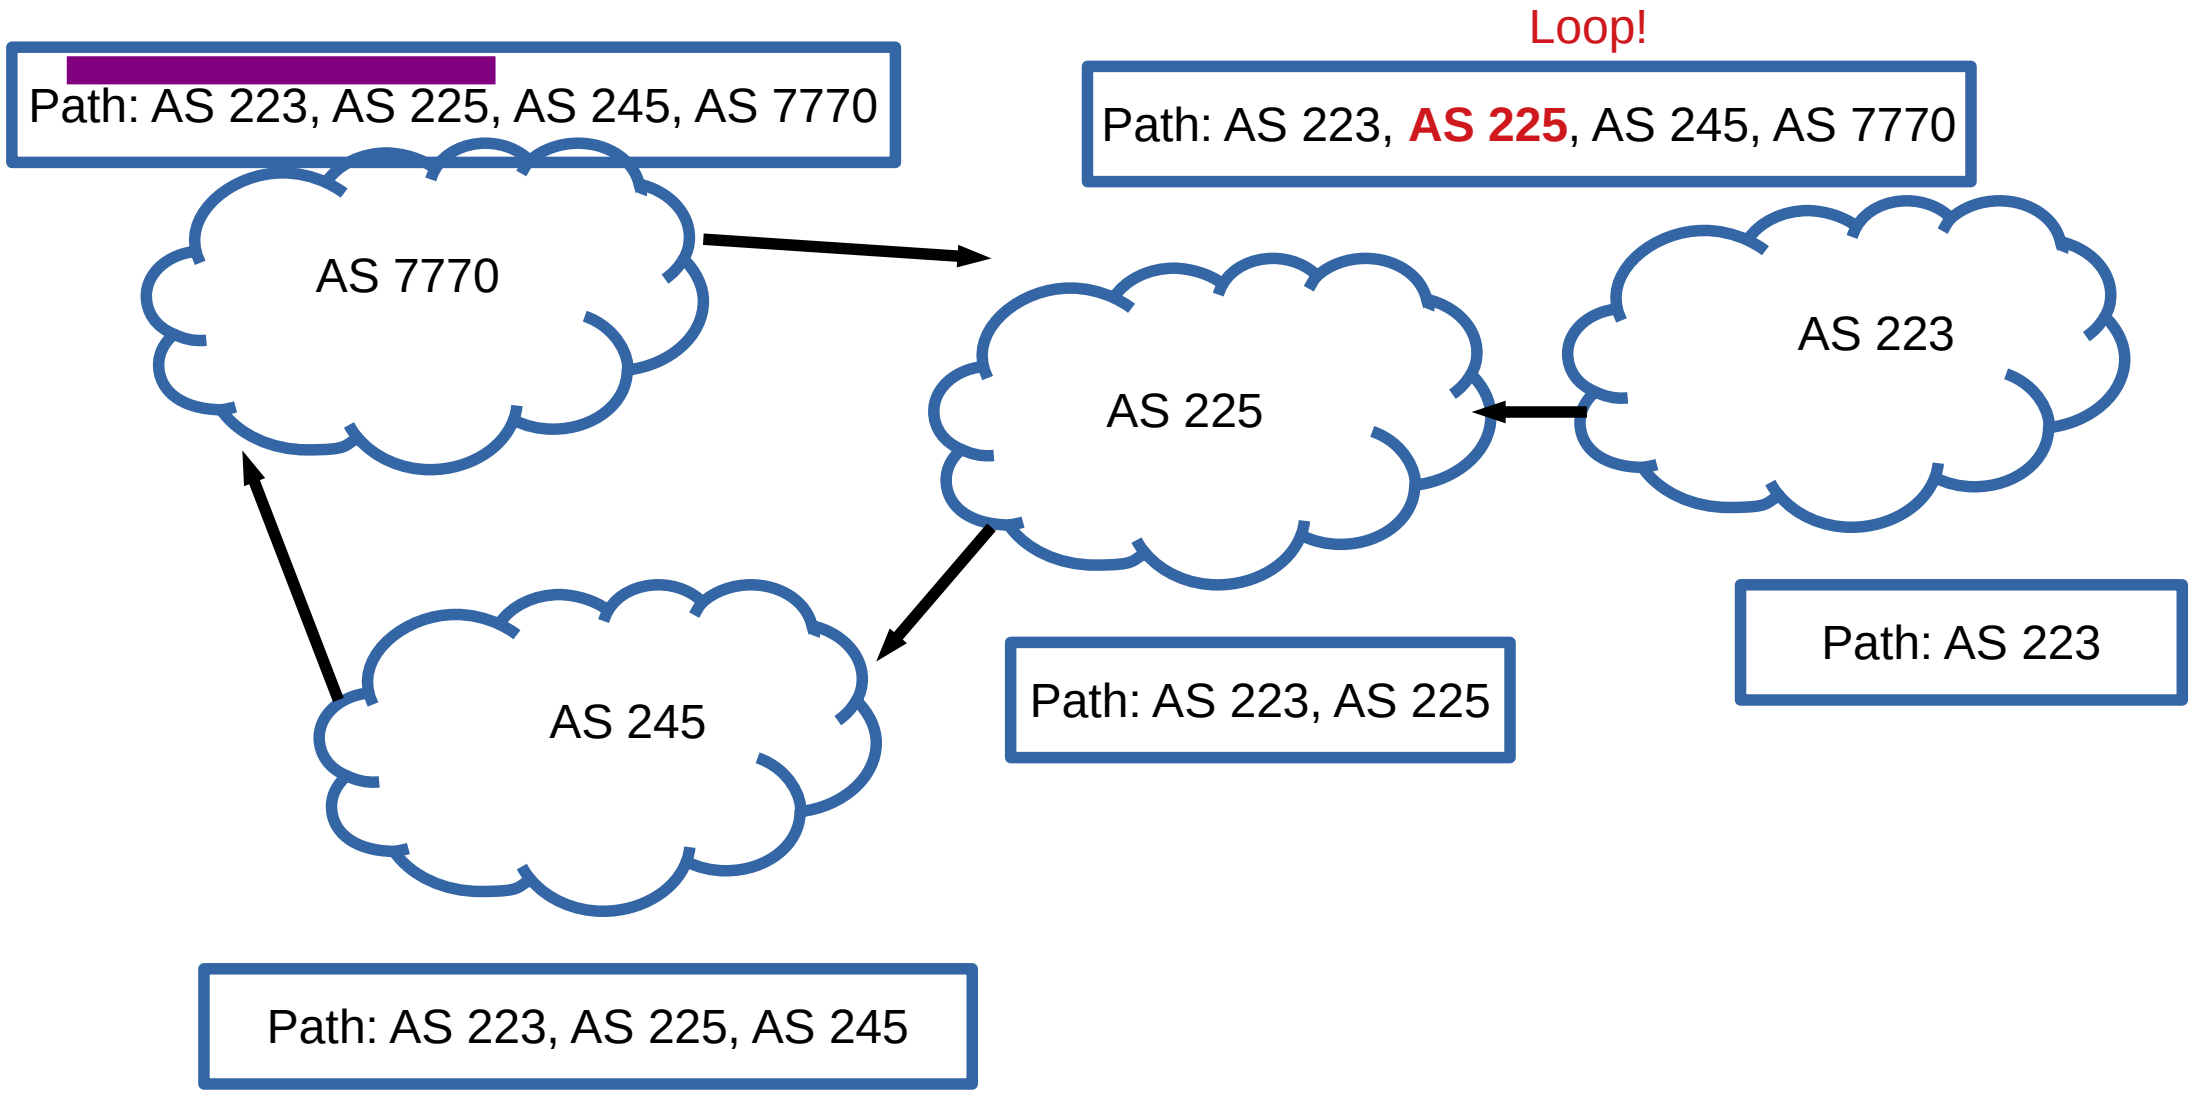

Transport layer

- So we have created a global network
- Two problems:
  - How do we turn this host-to-host packet delivery service into a process-to-process communication channel?
  - How do we look things up? ← More on this later

Transport services and protocols

- provide *logical communication* between app processes running on different hosts
- transport protocols run in end systems
  - send side: breaks app messages into *segments*, passes to network layer
  - rcv side: reassembles segments into messages, passes to app layer
- more than one transport protocol available to apps
  - Internet: TCP and UDP

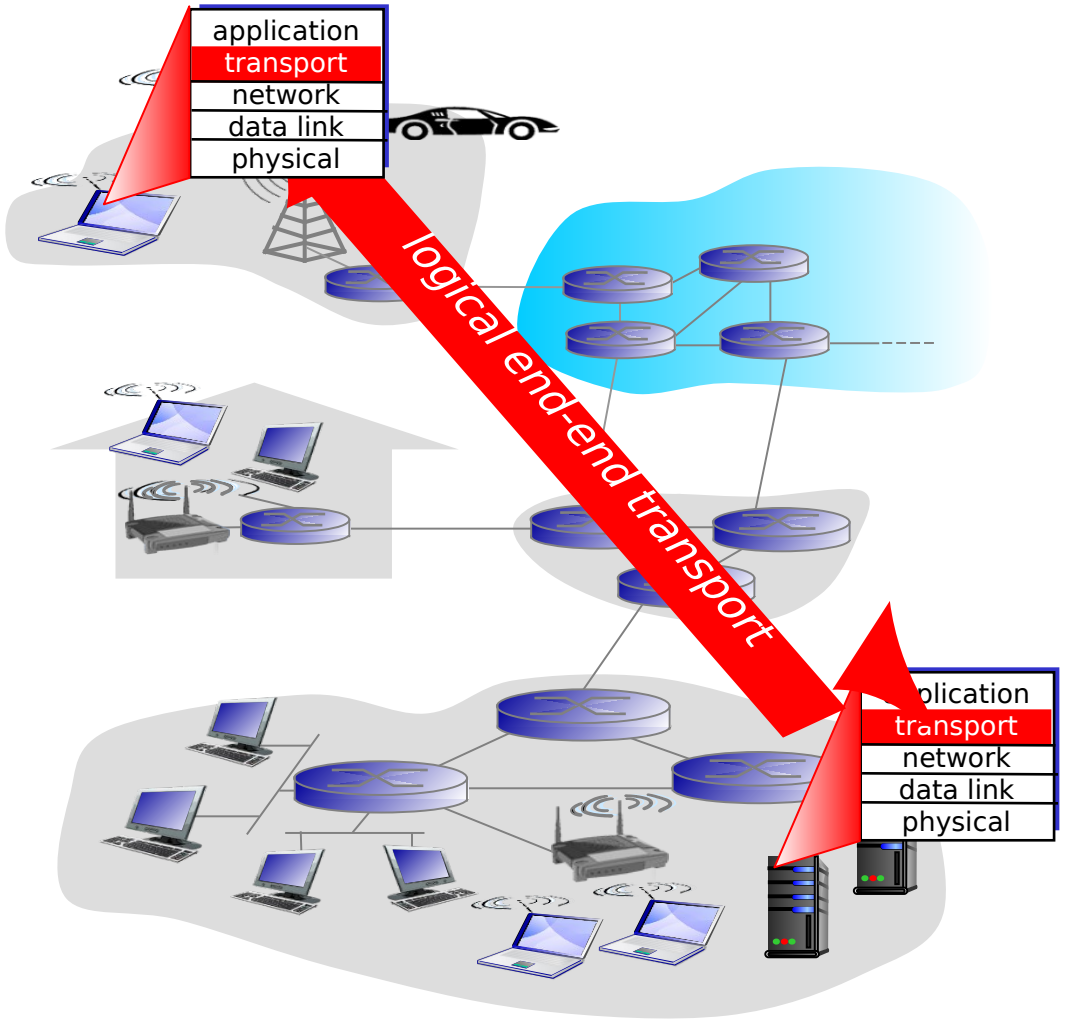

Multiplexing/demultiplexing

multiplexing at sender:
handle data from multiple sockets, add transport header (later used for demultiplexing)

demultiplexing at receiver:
use header info to deliver received segments to correct socket

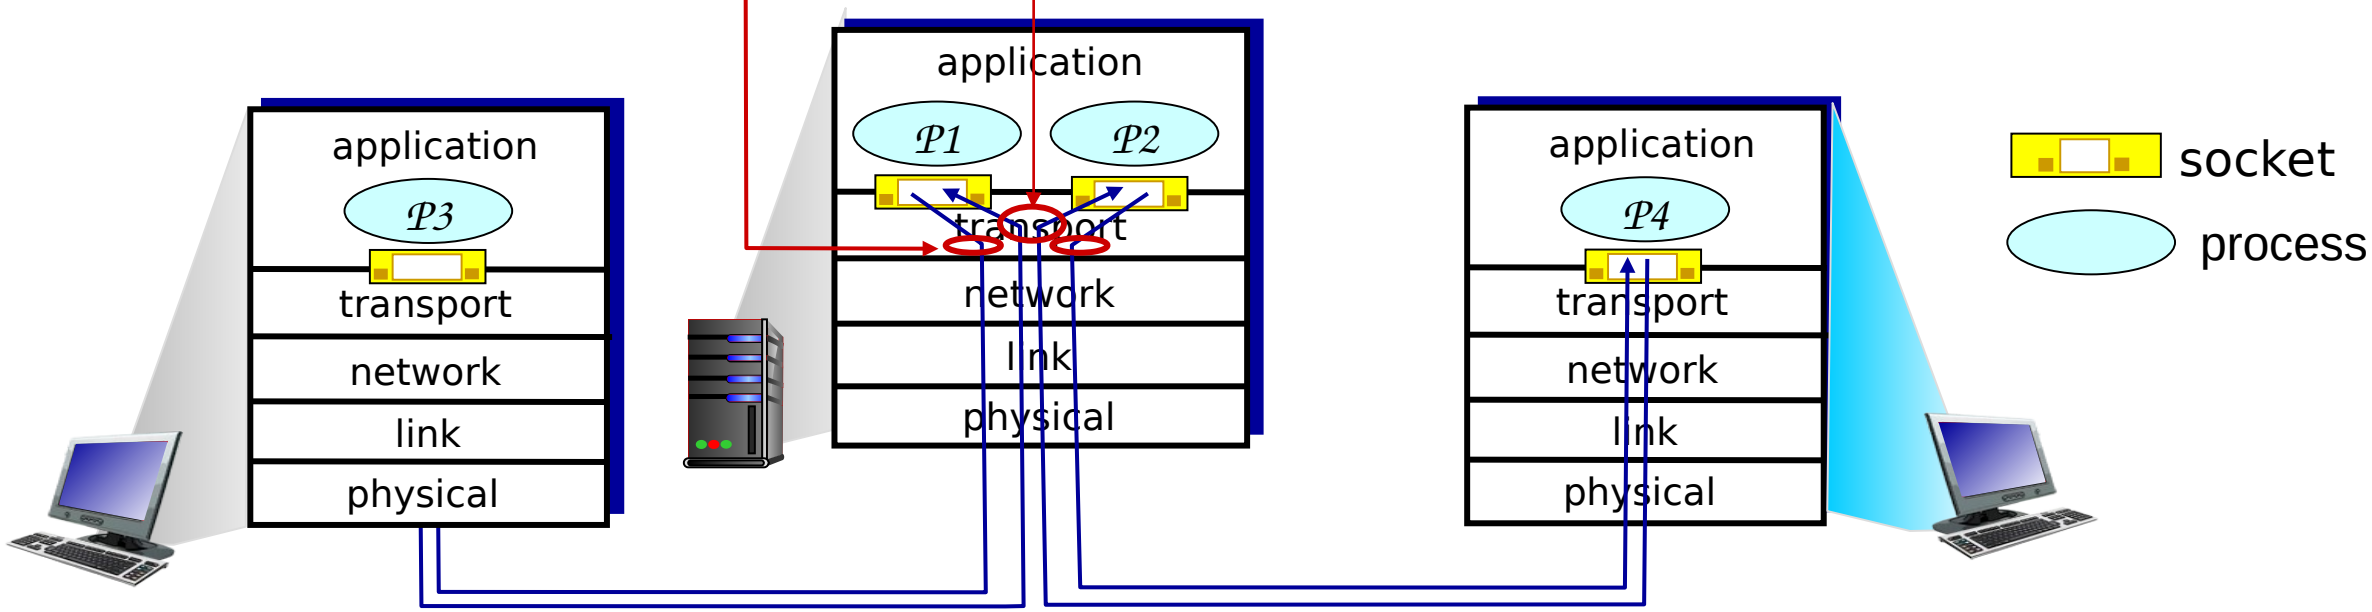

Finally – the DNS

- Internet’s telephone directory!

IP Based Communication

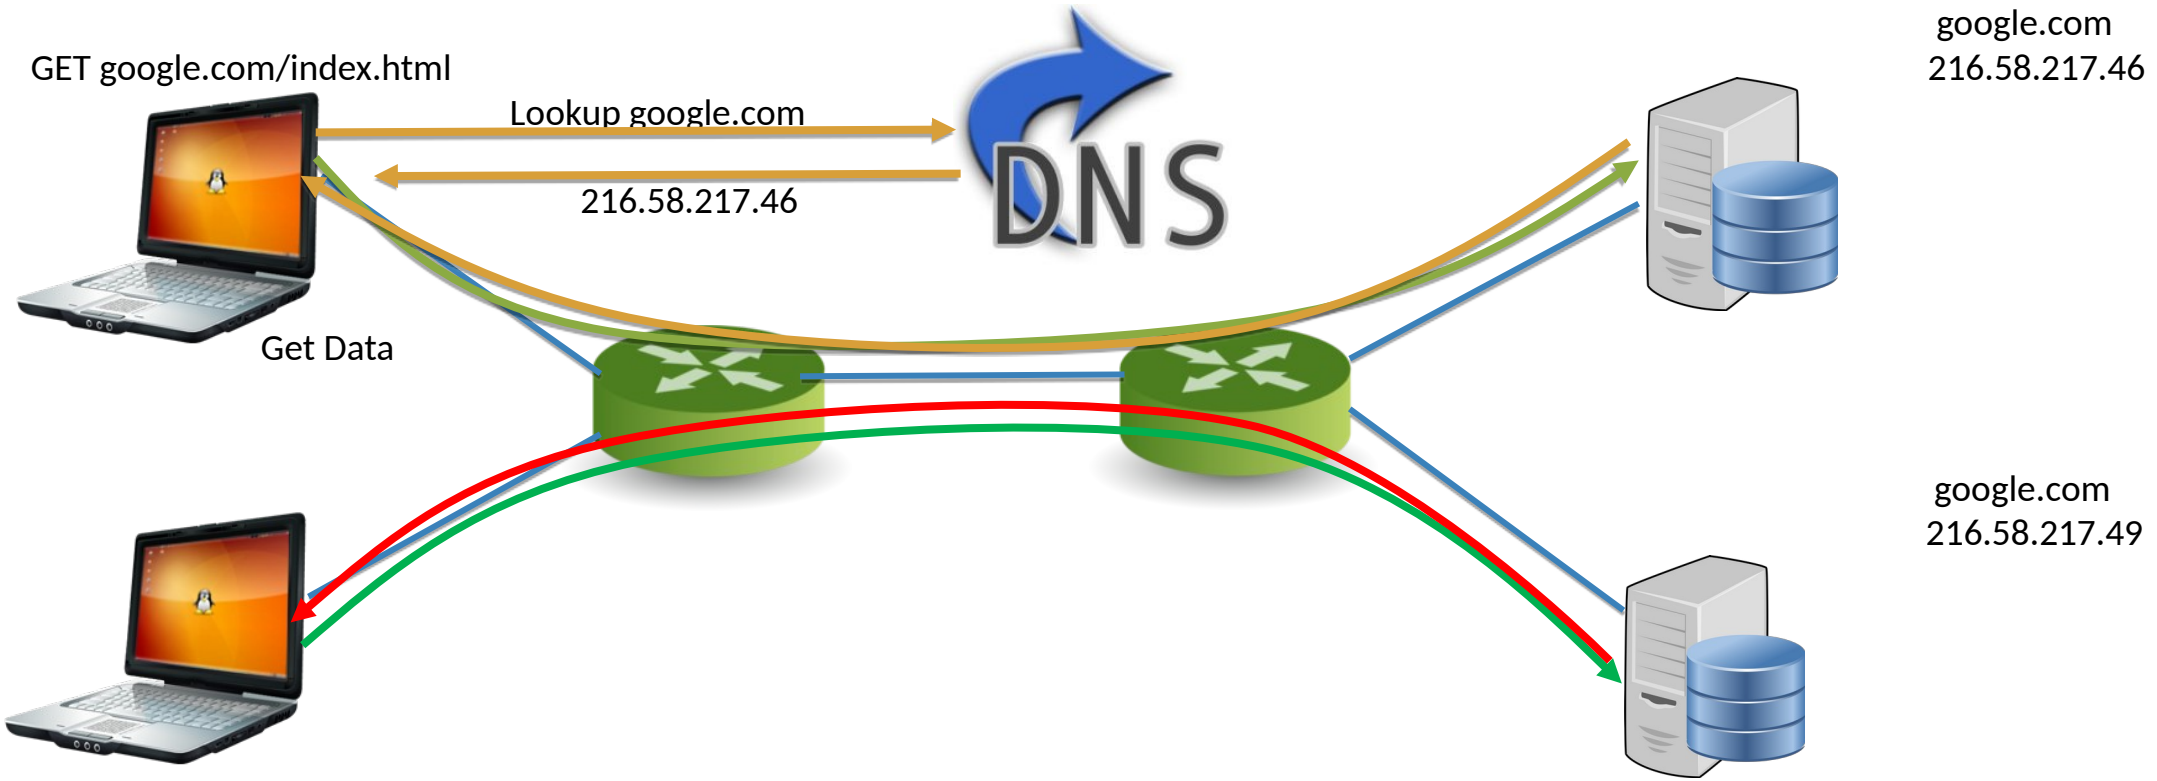

DNS – IP to Name

People: Good with names
Machines: Good with numbers

https://cat-bounce.com/ → 208.113.161.95

DNS maps IP addresses to human readable names.

DNS: a distributed, hierarchical database

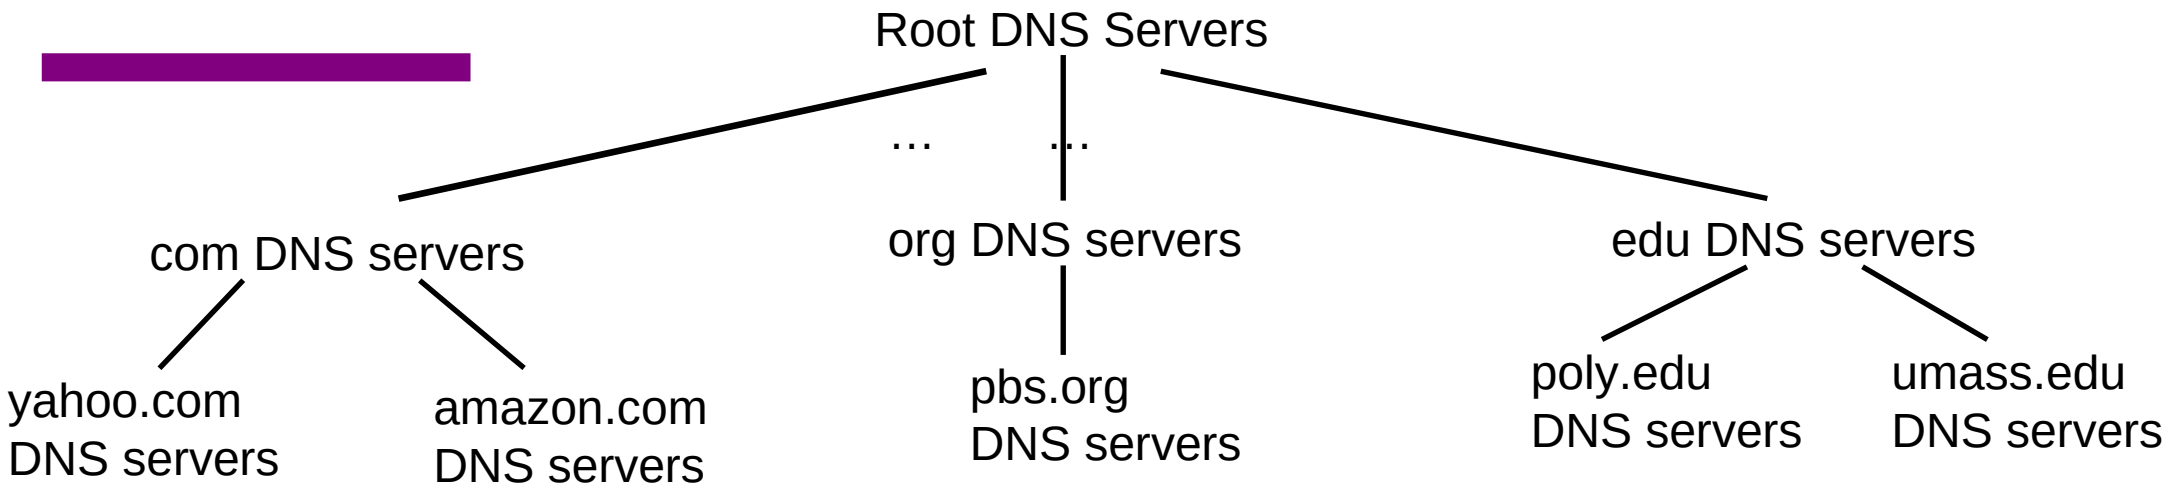

client wants IP for www.amazon.com;

1) client queries root server to find com DNS server
2) client queries .com DNS server to get amazon.com DNS server
3) client queries amazon.com DNS server to get IP address for www.amazon.com

DNS: root name servers

- contacted by local name server that can not resolve name
- root name server:
  - contacts authoritative name server if name mapping not known
  - gets mapping
  - returns mapping to local name server

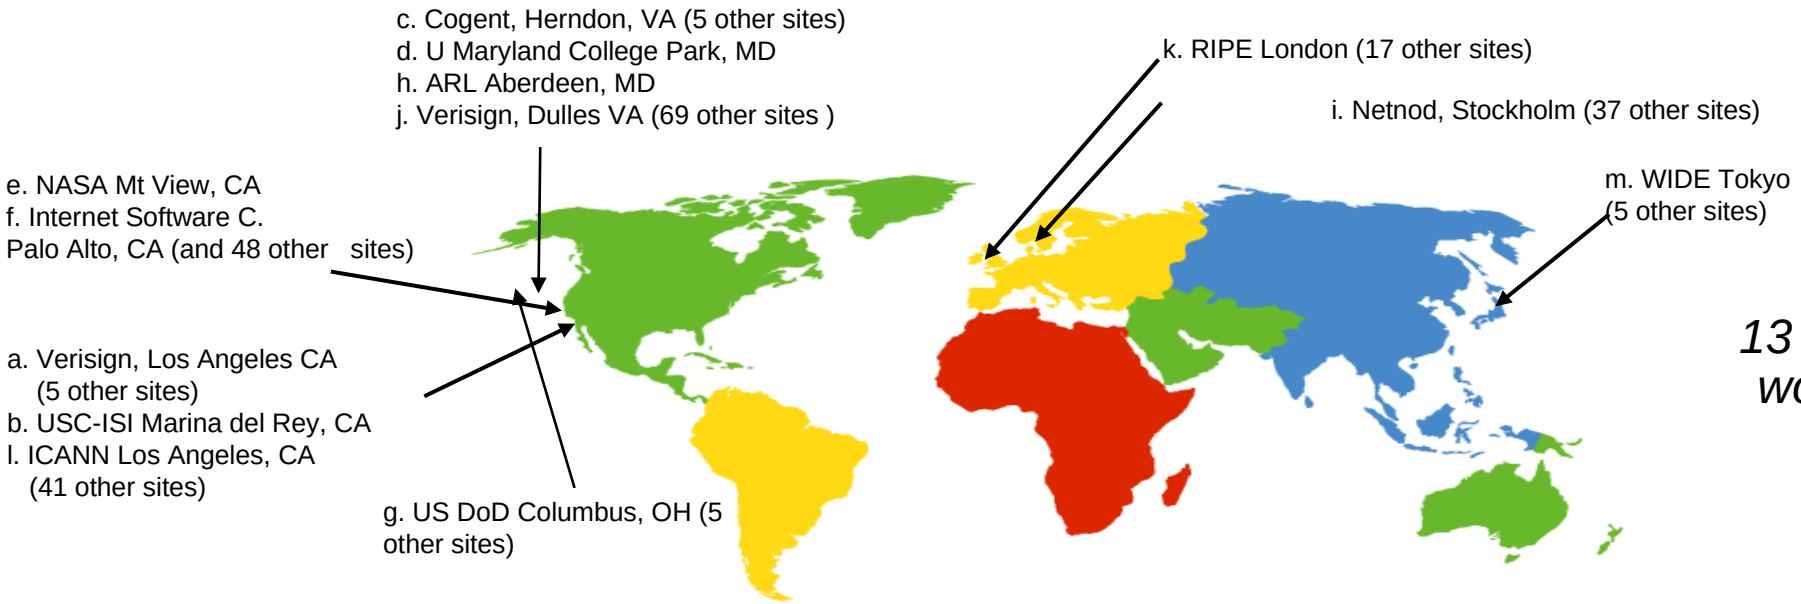

13 root name "servers" worldwide

Topics covered in this Lecture

- Where the Internet came from
- Design decisions
- Mechanisms that drive it
- Problems
- Where do we go from here?

Where do things break?

Link layer

– whale broke my fiber. Solar radiation. Your microwave.

- Also, TCP/IP is point to point – so you can’t do effective broadcast/multicast – even though your link layer is broadcast by nature.

Where do things break?

TCP/IP Layers

– Mobility!
- TCP pseudo header – connections are end-to-end!
- Multipath routing
- Content reuse

Where do things break?

Security

- No built-in security mechanisms
– DDoS
- Spam and Phishing
- Social engineering
- Insecure communication
- Password
- Human errors

DNS Poisoining

DDoS attacks

- Bombard root servers with traffic
- Bombard TLD servers
  - Potentially more dangerous
- Man-in-middle
  - Intercept queries
- DNS poisoning
  - Send bogus replies to DNS server, which caches

Redirect attacks

- Send queries with spoofed source address: target IP
- Requires amplification

BGP: Security problems

Anyone can advertise anything!!!

c|net BEST PRODUCTS REVIEWS NEWS VIDEO HOW TO SMART HOME CARS DEALS DOWNLOAD JOIN / SIGN IN

CULTURE

How Pakistan knocked YouTube offline (and how to make sure it never happens again)

YouTube becoming unreachable isn't the first time that Internet addresses were hijacked. But if it spurs interest in better security, it may be the last.

BY DECLAN MCCULLAGH | FEBRUARY 25, 2008 4:28 PM PST

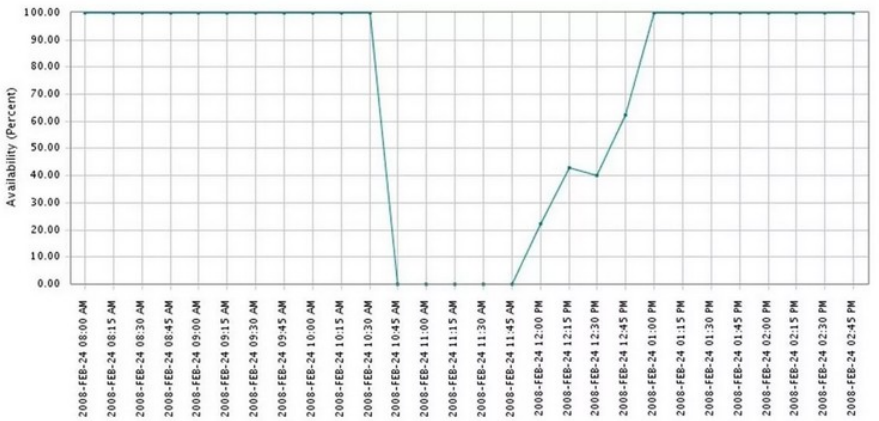

This graph that network-monitoring firm Keynote Systems provided to us shows the worldwide availability of YouTube.com dropping dramatically from 100 percent to 0 percent for over an hour. It

So What’s next?

How do we create protocols that interconnects new devices?

Clouds.
Edge Devices.
IoT.
Vehicular networks.

So What’s next?

The goal of this class – explore what I just presented in a lot more details.

- Explore how do we create protocols that interconnects new devices?

- CDNs.
- Clouds.
- Edge Devices.
- IoT.
- Vehicular networks.

References

A Brief History of the Internet - Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, Stephen Wolff
https://arxiv.org/html/cs/9901011

Vint Cerf: Re-Thinking the Internet -
https://www.youtube.com/watch?v=hagxPPoMGjw